

AD-A183 551

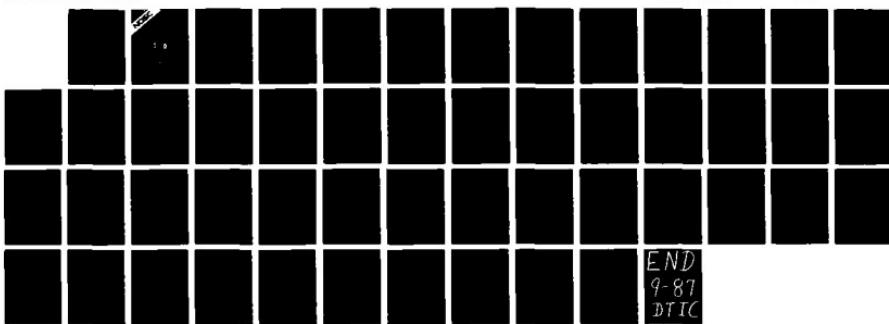
KNOWLEDGE ACQUISITION METHODOLOGY(U) SYSTEMS
EXPLORATION INC SAN DIEGO CA B D FRASER JUN 87
NOSC-TR-1094 N66001-86-D-0097

1/1

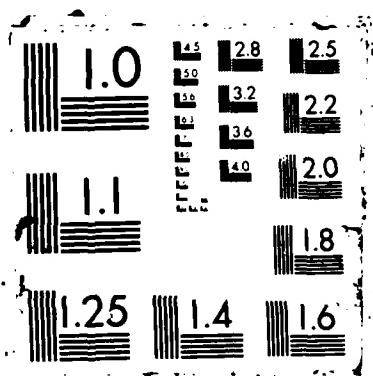
UNCLASSIFIED

F/G 5/2

NL



END
9-87
DTIC



DTIC FILE COPY

NOSC

NAVAL OCEAN SYSTEMS CENTER San Diego, California 92152-5000

12

AD-A183 551

Technical Report 1094
June 1987

Knowledge Acquisition Methodology

Bonnie D. Fraser



Approved for public release;
distribution is unlimited.

87 8 4 010

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

111
REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b RESTRICTIVE MARKINGS	
2a SECURITY CLASSIFICATION AUTHORITY		3 DISTRIBUTION/AVAILABILITY OF REPORT	
2b DECLASSIFICATION/DOWNGRADING SCHEDULE		Approved for public release; distribution is unlimited.	
4 PERFORMING ORGANIZATION REPORT NUMBER(S) NOSC Technical Report 1094		5 MONITORING ORGANIZATION REPORT NUMBER(S)	
6a NAME OF PERFORMING ORGANIZATION Systems Exploration, Inc.	6b OFFICE SYMBOL <i>(if applicable)</i>	7a NAME OF MONITORING ORGANIZATION Artificial Intelligence Technology Branch	
6c ADDRESS (City, State and ZIP Code) 4141 Jutland Drive San Diego, CA 92117		7b ADDRESS (City, State and ZIP Code) Naval Ocean Systems Center San Diego, CA 92152-5000	
8a NAME OF FUNDING/SPONSORING ORGANIZATION Office of Naval Technology	8b OFFICE SYMBOL <i>(if applicable)</i>	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N66001-86-D-0097	
10 SOURCE OF FUNDING NUMBERS		11 TITLE (Include Security Classification) Knowledge Acquisition Methodology	
PROGRAM ELEMENT NO 62757N	PROJECT NO CC96	TASK NO RS575	AGENCY ACCESSION NO DN188 532
12 PERSONAL AUTHORISI Bonnie D. Fraser			
13a TYPE OF REPORT Final	13b TIME COVERED FROM _____ TO _____	14 DATE OF REPORT (Year, Month, Day) June 1987	15 PAGE COUNT 50
16 SUPPLEMENTARY NOTATION			
17 COSATI CODES		18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number) artificial intelligence walk-through read-through expert system knowledge acquirer interview techniques	
19 ABSTRACT (Continue on reverse if necessary and identify by block number) The object of the work described in this report was to develop unobtrusive, reliable, and effective techniques for acquiring the knowledge needed to build the Airstrike Planning Advisor expert system.			
20 DISTRIBUTION AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21 ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a NAME OF RESPONSIBLE INDIVIDUAL R.H. Farrin		22b TELEPHONE (Include Area Code) (619) 225-7229	22c OFFICE SYMBOL Code 444

DD FORM 1473, 84 JAN

83 APR EDITION MAY BE USED UNTIL EXHAUSTED
ALL OTHER EDITIONS ARE OBSOLETE

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

EXECUTIVE SUMMARY

OBJECTIVE

Develop unobtrusive, reliable, and effective techniques for acquiring the knowledge needed to build the Airstrike Planning Advisor (ASPA) expert system. Provide guidelines and insights that can be applied to other Navy expert systems in the future.

RESULTS

A set of ideas about how knowledge acquisition is most effectively conducted was the primary result of the literature review and project experience.

RECOMMENDATIONS

Tailor methods to conform to two considerations: characteristics of people's ability to report what they know, and characteristics associated with the knowledge acquirer's ability to comprehend what is reported.

Accession For	✓
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification
By
Distribution	1
Availability Checks	
Dist	Approved by Sp. Ed.
A-1	



CONTENTS

	Page
INTRODUCTION	1
Problem.....	1
Objective.....	1
Approach.....	1
Results.....	1
RESULTS OF THE SURVEY OF ARTIFICIAL INTELLIGENCE LITERATURE.....	2
Problems Cited in AI Literature	2
Recommendations from AI Literature	2
Where to Start	3
PSYCHOLOGICAL AND FIELD-STUDY LITERATURE.....	5
Characteristics of People's Ability to Report.....	5
The Knowledge Acquirer's Ability to Comprehend.....	6
KNOWLEDGE ACQUISITION IN A NAVY OPERATIONAL SETTING.....	7
Best Expert Knowledge.....	8
Vocabulary.....	8
Finding Experts.....	8
Credibility	8
Specialization	9
Access	9
Setting.....	9
Unobtrusiveness	9
GATHERING KNOWLEDGE.....	10
Interviews.....	10
Questionnaires.....	15
Walk-Throughs	18
Read-Throughs	20
Observation	21
Recording the Data	23
Documentary Sources.....	23
FORMALIZATION OF KNOWLEDGE.....	24
KNOWLEDGE ACQUISITION CHECKLISTS	30
DISCUSSION	32
Acquiring Knowledge from the Expert	32
Controlling the Knowledge.....	32
Translating the Knowledge to Computer Compatible Form.....	33

	Page
RECOMMENDATIONS	34
Overview	34
Vocabulary Glossary	34
Procedures	34
Methods	35
REFERENCES.....	39
BIBLIOGRAPHY	43

ILLUSTRATIONS

1. Flowchart showing division of labor in a typical airstrike planning session.....25
2. Airstrike planning sequence27

TABLE

1. Knowledge acquisition methods36

INTRODUCTION

PROBLEM

Recent breakthroughs in computer technology have made it possible to develop systems which perform many of the functions normally done by human experts. These "expert" systems incorporate a great deal of expert knowledge, and rules for applying that knowledge. Expert knowledge is scarce, and knowledge acquisition is the bottleneck that most restricts application of expert system technology. Knowledge engineering--the process of collecting and refining the knowledge used in these systems--is expensive and time consuming. There are few guidelines for knowledge acquisition in general, and fewer for doing it in a specifically Navy setting.

OBJECTIVE

The principal objective of the work described in this report was to develop unobtrusive, reliable, and effective techniques for acquiring the knowledge needed to build the Airstrike Planning Advisor (ASPA) expert system. This system will develop technologies, methodology, principles, and standards for the use of expert systems in the Navy. The goal of this demonstration is to improve the timeliness and effectiveness of airstrike mission planning and decision making for a carrier wing. The technical issues are structured techniques for knowledge acquisition, knowledge validation, user-interface design guidelines and design tools, hybrid knowledge representation, and integrated decision aids. An additional long-term goal is to provide guidelines and insights that can be applied to other Navy expert systems.

APPROACH

Artificial intelligence (AI) and behavioral sciences literature was reviewed for problems, lessons, and recommendations related to knowledge acquisition. A number of knowledge acquisition techniques were attempted as an initial step toward evaluating their usefulness in overcoming those problems.

RESULTS

The literature research fell mainly into two categories. One was a review of the knowledge acquisition problem based on other artificial intelligence (AI) project experiences. The other category, psychological and field study literature, gave additional insights into the knowledge acquisition process. Combined with project experience, the primary result was a set of ideas about how knowledge acquisition is most effectively conducted. It was concluded that the most effective route entails tailoring methods to conform to two considerations: characteristics of people's ability to report what they know, and characteristics associated with the knowledge acquirer's ability to comprehend what is reported.

RESULTS OF THE SURVEY OF ARTIFICIAL-INTELLIGENCE LITERATURE

Very little discussion in AI literature is devoted specifically to knowledge acquisition, even though it is considered to be the main bottleneck in building expert systems. The survey of project lessons and issues by Welbank (1983) was a primary reference for this study.

PROBLEMS CITED IN AI LITERATURE

A frequently cited difficulty concerns the way experts express themselves. The problem is stated strongly by Barstow (1979): "In most domains, the ideas and structures are not known 'a priori' but must be discovered in a morass of knowledge that is available only informally or subconsciously." Experts may not be aware of all the knowledge they use, nor of their real problem solving strategies (Buchanan, 1979; Feigenbaum and McCorduck, 1983). Although both willing and intending to cooperate, they may give very poor descriptions of what they do. Or, what they have to say is ". . . not as explicit or precise as the knowledge acquirer would like" (Welbank, 1983).

Another problem cited in AI literature relates to interviewing, the knowledge acquisition method most often used. Welbank (1983) says there appears to be a risk of the knowledge acquirer taking over the interview, interrupting the expert and interpreting everything the expert says. Knowledge acquirers may provoke resentment by rejecting the expert's description of his reasoning and pressing him to justify every conclusion (Fox, 1983; Grover, 1983). In addition to alienating the expert, an overly aggressive interviewer may bias the knowledge by imposing a structure on it that does not actually correspond to the expert's reasoning. Similarly, an expert himself may impose an artificial and inaccurate structure on his own knowledge when he becomes acquainted with AI concepts, such as production rules, and expresses his thinking in those terms (Clancey, 1981; Smith and Baker, 1983). The resulting distortions lead to an expert system that has to be overhauled.

RECOMMENDATIONS FROM AI LITERATURE

A primary recommendation in knowledge acquisition is to ask specific questions. Construction of these questions is a major knowledge acquisition problem, since the knowledge acquirer does not know exactly what kind of knowledge is being sought (Welbank, 1983).

Supplying contexts for experts has been found to be a positive tactic for helping experts express what they know. "Experts may not be able to remember conditions, given actions, but they can remember actions given conditions" (Welbank, 1983). Here again, the difficulty is that the interviewer must first have sufficient knowledge of the domain to provide the expert with a context. Experience has shown that experts can criticize someone else's conclusions more readily than describe their own reasoning. Consequently, it is often possible to get experts to articulate their own reasoning by encouraging them to criticize the actions or plans of other experts (Waterman and Jenkins, 1979; Brooks, 1983).

Asking experts to describe a critical incident is another tactic. If the incident is especially memorable, the expert will be able to report details with a higher degree of accuracy than if he is asked to answer general questions about the domain. Because experts may not express their knowledge accurately or thoroughly, verbal accounts should be validated by observing their actual behavior (Waterman and Jenkins, 1979; Bainbridge, 1981; Fiegenbaum and McCorduck, 1983).

Some recommendations apply to the knowledge acquirer's preparation and conduct. The knowledge acquirer should know enough about the field to interpret what is said by the expert and to maintain credibility with the expert (Buchanan, 1979; Bainbridge, 1981; Hartley, 1982; Fiegenbaum et al., 1983; Grover, 1983). About a week of domain-related reading is described as typical. The knowledge acquirer should use the expert's vocabulary when communicating, not programmer's jargon. Technical jargon compounds opportunities for confused communication (Buchanan, 1979).

A general rule in the literature is to take a developing system back to the expert for repeated criticism. This method becomes a tool for knowledge acquisition, since experts are better at identifying exceptions to rules when they see something done incorrectly. Feedback from experts help keep programming efforts in tune with expert thinking (Gorry, 1973; Fiegenbaum and McCorduck, 1983; Fox, 1983).

WHERE TO START

There is a consensus in the AI literature that the first step in knowledge acquisition is understanding the structure of the domain's problem. There are no guidelines for how to do this, and what is meant by "understanding the structure" varies. Except for medical diagnosis, where much AI work has been done, classification of the structural elements of different kinds of domain problems remains to be done. For example, Chandrasekaran (1985) is pursuing the notion of "generic tasks," the idea that the kind of knowledge and control regime common to diagnostic reasoning is different from that common to, say, a design problem.

One view suggests that the initial knowledge acquisition effort should be devoted to assembling the domain's relevant concepts and attributes. Structure is found in the domain's objects, parts, subparts, and object attributes.

Clancey (1981) recommends first identifying the problem solving process characteristic of a domain, to produce a model of the reasoning strategy that is used. That model, in turn, will give the knowledge acquirer a good understanding of the problem structure and permit specific questions to be asked.

The rapid-prototyping school recommends building a system as quickly as possible based on a sample problem solution by an expert. They anticipate that the initial understanding of the problem will be deficient but that structure is implicit in the architecture of the prototype. Knowledge acquisition is conducted during iterative rounds of expert criticism of the system. The prototype is changed or rebuilt as more knowledge accumulates. Consequently, the first

major knowledge acquisition problem is to find a cooperative "star" expert with sufficient commitment and time to dedicate to the project (Buchanan, 1979).

A considerably different point of view is represented by Grover (1983). Maintaining that methods which closet the experts with the knowledge acquirer for months are not workable in the business world, he recommends that the first stage in knowledge acquisition should aim to produce a handbook containing the following: a general problem description; a bibliography of reference documents; a glossary of terms; a list of experts in the field; some reasonable performance metrics; and descriptions of typical reasoning scenarios.

Mittal and Dym (1985) argue that the first concern of a knowledge acquisition effort should be to conduct the preliminary research to ensure that the prospective system will be useful and used. Recommendations are made for accomplishing this. Multiple experts should be consulted for their conception of the domain problem, and specific attention should be given to identifying actual problems with which the expert community would like help. In order to take into account subdomain interdependences, an understanding of the whole problem should be achieved before a problem subdomain is selected. Attention should be given to ensuring experts really are experts. Consulting only one expert may produce a biased view of the problem because of his specialization. In contrast, observing how different experts solve a problem may highlight common essentials in their approaches.

A recent point of view maintains that a thorough analysis of the problem should be completed before system architectural decisions and implementation efforts proceed (deGreef and Breuker, 1985). The first knowledge acquisition step should be "... an analysis of the functions, the environment, and the users of the expertise, to arrive at the definition of the operational characteristics of the prospective system." Analyses of static domain knowledge, and then of the expertise in action, should follow. Experts evaluate the conceptual structure of a task, expressed in concept hierarchies. The focus is on the problem solving process, not system performance. A conceptual structure is considered easier to change than a prototype, and less obscure to an expert.

Various schools of thought maintain that there are different places from which the initial knowledge acquisition effort should start. Precisely how to capture the knowledge is not explained in detail. Nearly all of the perspectives begin with an expert, from whom the problem structure or definition is elicited. Hence, the common theme in the AI literature is the difficulty associated with eliciting that knowledge.

PSYCHOLOGICAL AND FIELD-STUDY LITERATURE

Experience with the ASPA project suggested that the knowledge acquisition bottleneck was not a single problem, but was, in fact, a result of two separate problems. One problem was the people's ability to report what they know or do, and the other problem was the knowledge acquirer's ability to comprehend what was reported or demonstrated.

CHARACTERISTICS OF PEOPLE'S ABILITY TO REPORT

Ericcson and Simons' (1980, 1984) model of human cognition and verbal reports provides insight into why people have difficulty reporting what they know. According to their model, those people who best qualify for the title "expert," because of their problem-solving short cuts, are also the ones most likely to have difficulty accurately reporting what they do. The short cuts, because of their utility in confining a computer's search space, are precisely the things a knowledge acquirer would like to capture. However, because of their routine nature, they may have slipped out of the expert's awareness.

Automatic Responses. According to many psychologists, routine actions or procedures quickly become automatic, bypassing conscious awareness. The kinds of procedures that move out of aware, cognitive control to automatic status include perceptual encoding (recognition), memory retrieval, and motor processes (e.g., riding a bike). The procedural knowledge is embedded in long-term memory, but it is difficult to access something that has not been explicitly memorized. A classic demonstration of this point involves asking people to report the number of windows where they live. A pause normally ensues while they envision the place, and then count the windows. Similarly, few people can reproduce the designs on common coins when asked, despite the number of times they have seen them.

The primary knowledge acquisition-related insight from literature on automatic responses is that people generally find it difficult to think in a vacuum, regardless of what they know. Retrieving information from long-term memory, which is anything that is not immediately engaging one's attention, is a painstaking activity. Consequently, as often observed during laboratory experimentation, simply asking someone to recount the steps they take in the course of their work will usually result in a gloss--an educated conjecture about what must be done. It is easier to reason than to remember. The well meaning answers resulting from simple, direct questions may not be sufficiently accurate or precise for a computer database.

Protocols. More accurate reporting can be obtained when it accompanies the actual doing of a task, at least if the reporter is properly instructed. Overly general questioning is an invitation to theorize rather than to report. Ericcson and Simon (1984) maintain that people are more capable of accurately describing the contents of their attention during or shortly after doing a task, although some people are better at it than others, and a heavy cognitive load may result in some information not being reported. The task itself provides cues, exercising the retrieval of information from long-term memory.

Recognition. Recognition is the basis of the easiest and most accurate reporting. We all do best on objective tests where possible answers are presented for our selection. If the material is familiar, such as a dollar bill, recognition is instantaneous. There is no laboring to retrieve information from long-term memory, because long-term memory is automatically triggered.

THE KNOWLEDGE ACQUIRER'S ABILITY TO COMPREHEND

It is difficult for the knowledge acquirer to achieve the competence required to elicit, and meaningfully interpret, the knowledge that experts convey. It is a chicken and egg dilemma. The obvious solutions--directly observing experts' problem solving activity or asking direct questions--beg the question. Seeing a process does not tell you what someone is thinking. The experts may be incapable of articulating their reasoning without considerable probing by an informed inquirer (Ericsson and Simon, 1984). Similarly, specific questions are best suited to acquiring accurate knowledge from an expert, by providing memory cues and minimizing the expert's chance to misunderstand the question. But a specific question assumes that the questioner knows a considerable amount. It also drastically restricts the range of possible answers (Flammer, 1981).

Levels of Questioning. There is some experimental evidence that there is a structure to comprehension, a hierarchical organization of knowledge, that influences the kind of questions a questioner will normally pose to achieve understanding (Flammer, 1981; Flammer et al., 1981; Flammer et al., 1984). There is also some evidence that both the number and quality of naturally occurring questions are related to what one already knows. Both points suggest where, from the knowledge acquirer's point of view, it is most efficient to start knowledge acquisition.

An earlier model for understanding stories was a good example of the hierarchical structure of knowledge. The vocabulary, the lowest level in the hierarchy, was followed by sentences, topic, author's message, and the context of the author's message. Experiments conducted in conjunction with this theory found that increased disturbance of a text, correlated with increased macrolevel questions, oriented toward understanding the gist of the whole text. The unanticipated finding was that microlevel questions, oriented toward understanding vocabulary, decreased with disturbances of the text. It is as though the subject has only so much attention to devote, and concentration on overall meaning preempted attention to more finely grained knowledge gaps. It appeared that "different levels of knowledge hierarchy also determine the priority of respective (knowledge) 'holes'" (Flammer, 1981).

Another study examined the quality and quantity of questions that were asked during the learning of new material. An easy, introductory manual with examples, and a more difficult, abstractly framed text were used in the experiments. One set of students had a fairly advanced background in the topic while the students in the other set were complete novices. The study found that the ability to ask questions was associated with how the groups' prior knowledge matched the material they were given. Advanced students formulated more questions and hypotheses when given the more difficult text; novice students did the same with the easier text. Conversely, when their level of knowledge was not matched to the text, both sets of students were less capable of formulating questions at all (Miyake and Norman, 1979). One can know too

much to have questions occur. More importantly, one can know too little to know where to begin to make sense of material.

Lacking both vocabulary and a basic understanding of fundamental concepts compounds problems of comprehension, especially in terms of retaining what is heard. Many of the basic concepts are in the spoken vocabulary. Clarification questions tend not to be asked, because attention is devoted to following what's going on (Flammer et al., 1984) and questions that are asked can be annoying or distracting to the speaker. Sharing English as a language may create the impression that real communication has taken place, because familiar rules of syntax govern the verbal report. But the momentary enlightenment that accompanies active listening without sufficient background is soon extinguished. "Although thoughts without words are possible, such thoughts are elusive and easy to forget" (Sowa, 1984).

When starting knowledge acquisition, the knowledge acquirer often will be ignorant of the esoteric language, objects, relevant concepts, and relations which distinguish a domain as an expert field. Hence, although the ultimate goal of knowledge acquisition is to elicit the expert's "rules of thumb," practices and thinking, that kind of knowledge may not be very useful to the knowledge acquirer in the early stages.

From the point of view of the knowledge acquirer, someone who is not so accomplished may be a better original source. Although introductory textbooks may present a naive version of how problems are really thought about and solved by experts in a domain, it is a version on which a better understanding can be built. "To ask about a concept, the subject must know what is missing and what is necessary for further understanding" (Miyake and Norman, 1979).

KNOWLEDGE ACQUISITION IN A NAVY OPERATIONAL SETTING

A major goal of the ASPA knowledge acquisition effort is to develop knowledge acquisition methods tailored to the special requirements of a Navy operational setting. Because those circumstances vary so much from the laboratory setting in which most expert systems have been developed, experience has been the primary guide for achieving that goal. The knowledge acquisition methods used in a Navy setting have to be evaluated by their unobtrusiveness, as well as by their ability to capture knowledge.

Lehner et al. (1985) specified four ways in which traditional expert systems differ from military expert systems. Three of these differences have a direct impact on knowledge acquisition.

1. The traditional systems addressed problem domains with a well established, well-documented, and static knowledge base. Military applications tend to involve ill-specified knowledge bases, where human experts differ considerably in their opinions.
2. In traditional systems, it was sufficient to model the system after one good human expert. In military applications, the system must often merge the expertise of a number of human experts with different areas of expertise.

3. In the traditional system, the assumed user community was not very diverse. Users of medical-diagnosis programs were likely to have some type of medical degree.... In many military applications, on the other hand, the level and type of experience of users will vary considerably.

These characteristics and others had a substantial influence on the ASPA knowledge acquisition effort.

BEST EXPERT KNOWLEDGE

Lack of documentation of basic expert knowledge, vocabulary, and fundamental concepts meant there was little opportunity for the knowledge acquisition team to study the field before approaching the experts. Documents and manuals were of little help, as can be seen from a team member's notes: "Textual materials are not organized in task descriptive or task analytic format.... Materials read like a potpourri of factual knowledge with control knowledge coverage non-existent." There are no introductory textbooks which teach airstrike planning. Schools teach aviators to fly, but the basics of the planning are learned through experience and on cruises. Advanced courses dedicated to refining skills are available, but they assume that attendees know the fundamentals.

VOCABULARY

Military vocabulary is highly specialized. A single object or concept may have a variety of names, nicknames, and acronyms. Members of the expert community themselves have difficulty keeping up with all of the acronym's meanings, and have their own incomplete dictionaries. Not only is the aviators' own professional jargon a formidable communication barrier, they are not comfortable with AI jargon. For instance, aviators and bombardiers were uncomfortable when referred to as experts, based on their belief that there are no real "experts" but, rather, there are those who are well respected and accomplished, based largely on experience.

FINDING EXPERTS

Finding credible experts, and gaining access to them, was a major issue in this setting. The first problem was to locate those who knew about the airstrike community. The next problem was to identify the established authorities, their locations, and their reputations.

CREDIBILITY

A guiding assumption of the knowledge acquisition effort was that ASPA's credibility would ultimately be determined by both the quality of knowledge it contained and the credibility of the contributing experts. Other military projects have had unfortunate experiences. Their systems were not used, because they had not taken into account the characteristics of the user community.

The overall community of naval aviators is relatively small, with only about 5,500 members. Within that community informal lines of communication and individuals' reputations are important. Finding well-respected experts was, therefore, a primary concern. That concern was vindicated on occasions when it was pointed out to the ASPA team that we should be wary of recruiting retired personnel. Although apparently desirable candidates because of their availability, they might be retired because they were not well thought of in their careers. Active-duty commanding officers and executive officers stressed that they should be the ones to select the contributing personnel to ensure their capability.

SPECIALIZATION

The other sense in which identifying credible expertise was important reflected the nature of the airstrike planning problem. Airstrike planning involves a number of specializations; no one person can be expert in all of them. A community of expertise was required. Identifying who would qualify, and where they were physically located, followed.

ACCESS

Although the aviation community is small, the Navy itself is a highly hierarchical bureaucracy. Obtaining necessary authorizations to gain formal access to the best experts caused considerable project delays, while the project was sold to officials who were concerned with its material benefits to the Navy. Operational people were favorably impressed that they were being enlisted as coparticipants in the system's development.

SETTING

Candidate experts were very willing to cooperate, but their time was not their own. They were not on the project's payroll and there was no career incentive for dedicating much of their time. In fact, those whose expertise was most sought after were those whose time was most committed.

UNOBTRUSIVENESS

Doing unobtrusive knowledge acquisition required going to where the experts were, and becoming accustomed to their environment. Experts have their own work to do, which calls for flexibility on the knowledge acquirer's part. In light of the context, methods which minimally interfere with daily work and careers are the most efficient. The values associated with the setting also affected methods. For instance, a highly secret environment sometimes made tape recording, and even note taking, problematic. It was difficult, for example, to gain access to some classes.

GATHERING KNOWLEDGE

A central issue in developing a systematic knowledge acquisition methodology is to identify techniques for defining a domain's problem(s) that do not alienate or exhaust experts, but which do succeed in breaking the problem into parts that are cognitively accessible to the knowledge acquirer. This involves tailoring methods to conform to people's strengths and shortcomings as accurate reporters. It also entails directing the goals of knowledge acquisition along the most efficient route to knowledge comprehension. Accumulated knowledge is the basis for progressively structuring knowledge acquisition methods. Hence, the knowledge acquirer's learning curve is a central key to efficient knowledge acquisition. The faster the knowledge acquirer advances along that curve, the faster he is able to structure knowledge acquisition methods. The structured methods are tools which enable experts to express what they know more accurately and completely.

A number of different approaches, methods, and techniques were used by the ASPA team, and are listed below with brief descriptions. Essentially, they are variations on four main themes: interviews, questionnaires, observation, and document searches. Knowledge sources included course materials, flow charts, and articles.

- Unstructured interviews. Neither questions nor answers are specified; the knowledge acquirer records the data.
- Open-ended interviews. Questions are specified, but answers are not; the knowledge acquirer records the data.
- Short-answer questionnaires. Questions are specified, answers are not; the expert records the data.
- Forced-answer questionnaires. Questions are specified, answers are specified; the expert selects.
- Walk-throughs. The task structures the data, which the expert then articulates; the knowledge acquirer records the data.
- Read-throughs. The document provides the expert cues. The knowledge acquirer asks questions. The knowledge acquirer records the expert's answers.
- Observation. Varies in the degree to which it is structured by questions. The expert performs the task and the knowledge acquirer records the data.

INTERVIEWS

A number of characteristics influence the degree to which the interviews should be structured (i.e., the degree to which the interviewer should be directive), their length, and where they should take place. There are also standards which are generally accepted conventions in

disciplines where the question of how to interview successfully is a major question. To elaborate on all of them would result in a different paper. However, in view of their relevance to knowledge acquisition interviewing, they will be highlighted here.

Those who are interested in interviewing have questioned what the interviewee's motivation is to participate in an interview. It is generally accepted that under conditions of trust, people find it rewarding to share what they know well; it satisfies altruistic, emotional, and intellectual interests (Hyman et al., 1954; Rogers, 1959). Trust can be affected by a number of issues; for instance, whether what is said will be on or off the record, what the information will be used for, whether the respondent will be judged for what is said, etc. It may take time to establish the requisite trust.

Listening is another important aspect of good interviewing. An inquiring stance conveys to the interviewee that what is said is valuable. Debates and interruptions undermine the idea that one is listening; the evidence is that the interviewer is preoccupied with composing a response. Studies have shown that repeated interruptions tend to extinguish the interrupted parties' participation in an exchange; they cease to talk (Natale et al., 1979; Rogers and Jones, 1975; Sacks et al., 1974; Zimmerman and West, 1975). Privacy is central to successful interviewing, and interviewer neutrality should be taken on issues that tend to divide groups. Status differentials should be minimized, who the interviewer is (the objectives of the interview) should be clear, and the interviewer should know to whom one is speaking.

Knowledge acquisition interviews are in a category by themselves. Often the questioner does not yet know enough, or is not working with an explicit hypothesis, and is, therefore, ill-equipped to structure the interview (that is, to anticipate the range of possible answers to the questions). Furthermore, the expert will often attempt to describe ideas that have never been explicitly articulated. Knowledge acquirers themselves can help or hinder the expert express what may be so well known that the expert may not be aware of it. Nondirective interview techniques which support the expert's conceptualization process, i.e., interviews, where the interviewer listens carefully and feeds back what is said, encourage experts to formulate and express ideas. The nature of the interviewing also affects its optimum length. The concentration that is required of both parties to the interview is demanding. Most experienced researchers and clinicians recommend an hour to an hour and one-half per interview. However, this is not always feasible in a field setting where one has often to go longer to take advantage of expert availability. In those circumstances, breaks are important. The optimum context is one which enables the interviewer to control the situation, minimizing interruptions and distractions. For example, a quiet office where phone calls and visits can be avoided is good.

Unstructured Interviews. In an unstructured interview, often called an exploratory interview, neither the specific questions nor the range of possible answers is anticipated. Instead, the answer to one question leads to the next question. In the earliest knowledge acquisition stage, exploring the domain, there is little alternative to conducting unstructured interviews, especially when the domain lacks a well documented knowledge base.

Unstructured interviews permit the expert, rather than the interviewer, to introduce concepts and vocabulary, giving the knowledge acquirer an initial sense of the domain. While these interviews can be frustrating, because so much is heard that makes so little sense, they do resemble the "total immersion" that is recommended for learning new languages. Unique language is an important distinctive characteristic of an expert domain.

Unstructured interviews were used on a number of occasions by members of the ASPA knowledge acquisition team. Their value was influenced by a number of factors. One factor was the knowledge acquisition team's degree of ignorance. When little was known about the domain, the unstructured interviews were not focused by a topic for fear of stifling the introduction of other important topics and information. The lack of structure, and apparent lack of "rigor" in the questioning, was trying to the members of the computer science team. They found the loosely organized material difficult to incorporate in a prototype system. Secondly, the size and complexity of a topic affected the unstructured interviews. When the problem was relatively small, like fusing, an introductory lecture was fairly enlightening and exhaustive. Individual experts varied in the way they conveyed their knowledge. Some experts' responses were fairly structured, despite the lack of structure to the interview. Finally, the characteristics of the knowledge acquirer had an impact on the evaluation of unstructured interviewing.

The problems associated with attempts to get an initial bearing on the airstrike planning problem, via unstructured interviews, were succinctly summed up in a team member's notes:

"The interview process was frustrating. The expert talked too fast and it was difficult to direct conversations given that our goals were nebulous (learn about strike planning). We received a data dump from the expert that was too detailed for some tasks (we couldn't write or take notes fast enough, or remember the details later on) or too general (the expert had a hard time conveying his control knowledge). Attempts to slow down the expert to cover vocabulary problems or explain concepts often led the interview astray into tangents where the original conversation flow was never recovered."

Interviewing as an Art. Some people maintain that knowledge acquisition is an art, since many of the skills of a good knowledge acquirer are not easily described or learned. For instance, being able to conduct unstructured interviews comfortably calls for a conviction that the outcome will be of value and an ability to tolerate a high level of ambiguity. The first requirement is a result of training. The ability to deal with ambiguity is more a matter of personality.

Harrison and Bramson (1982) maintained that there are five styles of thinking which, as ideal types, are suggestive of the ways personality can be important. These styles of thinking are synthetic, idealistic, pragmatic, analytic, and realistic. People vary in the degrees to which their thinking favors one or more of these styles.

- Synthetic thinkers tend to focus on abstract concepts, and consider data meaningless without interpretation.

- Idealistic thinkers are interested in human values and processes; theory and data are given equal value.
- Pragmatists are eclectic, interested in any theory or data that works.
- Analytic thinkers favor formal logic and deduction, and value theory and method over data.
- Realists are empiricists who rely on facts and induction, and value data over theory.

Each style of thinking is at its best in different kinds of situations, e.g., a strong realist functions best in well defined, objective situations, but may rush to oversimplified conclusions on the basis of perceived "facts." A combined synthetic/pragmatic reasoner is said to be best equipped to sustain ambiguity.

Not surprisingly, there is some correlation between people's styles of thinking and their occupations. Considering that a strong conceptual thinker, a synthesist, makes a fine computer scientist, but also tends to seek conflict, it is also no surprise that he can appear argumentative in other circumstances. A strong realist, interested in facts, produces a robust computer system, but may feel completely adrift in an unstructured situation.

The point here is not to make an argument that biology is destiny, but rather to make the obvious, though often neglected, observation that people differ in their interests and in what "comes naturally" to them, i.e., the questions they will pose in a new situation, where they will feel most comfortable, and for what they are best suited. A major rationale for any methodology is to diminish the haphazardness associated with letting different natural aptitudes gravitate to what is most comfortable; the goal of a knowledge acquisition methodology is to place the process on a steady, systematic, and objective footing.

Open-Ended Interviews. Open-ended interviews introduce some structure into the knowledge acquisition process. The object of introducing structure is to better control the information, not the expert. Questions are specified by the interviewer, but responses are neither anticipated nor standardized. Open-ended interviews enable the interviewer to direct the knowledge acquisition process in two important ways: determining the level of questioning and providing a focus that makes expert digressions more tolerable. Digressions are not only often highly informative, they are a way that an expert assures himself that important information, perhaps with life and death implications, has been conveyed. The pre-set questions refocus the interview after a digression.

Open-ended interviews presume that the knowledge acquirer is more or less aware of the kind of knowledge he is after. Although the level of questioning does not necessarily produce the same level of response, general questioning can elicit various frames of reference and highly specific questions can elicit factual and technical detail. But often, regardless of the level intended by the questioner, highly detailed responses will result, as the expert provides examples

while working through the question or an issue that has occurred to him. In such a case, the knowledge acquirer must extrapolate the "direct" answer implicit in the examples.

Open-ended interviews are frequently used by members of the ASPA team. Samples of general questions are:

- What do you think is the most important part of weaponeering?
- What do you think is the hardest task?
- Given your knowledge of computer programs, why or why not would weaponeering be difficult for a computer to do?
- Very briefly, can you list the main tasks associated with weaponeering?

Examples of specific questions are:

- How are weapons and fuel tanks usually configured on an A-6?
- On p. 14-11, what does "AERO 7A," "AERO 7B" mean?
- Are external fuel tanks dumped before the target area?
- Is it much trouble to switch fuel configurations?

Two of the problems with open-ended interviewing are that important issues may not be anticipated by the questions, or the questions themselves may miss the point. Adhering too rigidly to a script can preclude access to information of equal or greater importance, and can annoy an expert. Consequently, it is important to remain responsive to what is being said and to the expert's nonverbal reactions. Additionally, direct questions may elicit a gloss by an expert, i.e., an answer which is not thorough or precise.

Open-ended interviews are, nevertheless, an efficient route to acquiring a first approximation of an expert's reasoning process. In order to minimize the problems, open-ended interviews are conducted, whenever possible, in conjunction with situations or memory aids for stimulating recall. For example, shortly after solving a weaponeering problem an expert was asked what weapon he selected and what factors drove his selection of that weapon. He easily named five: (1) it was a low threat target; (2) it was a point target; (3) the weapon was precise; (4) the weapon was plentiful; and (5) the weapon was inexpensive.

Group Interviews. Group interviews can be conducted in either open-ended or structured response formats. The points which recommend group interviews as a knowledge acquisition method are stated well by Thompson and Demerath (1952). "Comments by one member of the group remind others of their experiences or of additional details. Furthermore,

since no individual need focus his entire attention on the details of the interviewer's next question, group members have more time for reflection."

The ASPA's knowledge acquisition team conducted two group interviews. The first was a highly structured interview, where both the questions and the possible answers were specified. The allowable responses were restricted to "always," "never" or "sometimes." A "sometimes" response required elaboration. The questions were designed to validate system rules and a portion of a task analysis. The questions were, therefore, specifically directed toward tapping the experts' long-term memory, and drawing out exceptions. The experts not only provided memory cues to each other during the interview but also criticized each other's responses and exceptions.

An entirely different kind of group interview was conducted in order to capture the reasoning associated with producing an airstrike plan. Exploratory in nature, an unstructured interview was focused by the experts' recent experiences in a planning session. The experts were asked to tell how they arrived at the particular plan chosen for that scenario. The goal was to reconstruct the reasoning associated with all of the important plan elements: weaponeering, tactics, route planning, support, electronic warfare, etc.

Members of the ASPA team who were present taped the planning session. The group interview, conducted fresh on the heels of the session, gave the experts the opportunity to reflect, discuss, and reconstruct their approach. They were able to name and explain the factors that had prompted, and those that would have changed, their decisions.

A group interview does not permit a knowledge acquirer to probe too deeply, except when knowledge has already been collected and can be incorporated in highly structured questions. There are other practical limitations to the procedure. Arranging for a group interview may be difficult. It also may be difficult to get the mix of experts right, so that one or more participants do not dominate or inhibit others (Thompson and Demerath, 1952).

QUESTIONNAIRES

Ideally, a skilled interviewer can help an expert express what he knows by encouraging him to relax and do the cognitive work necessary to formulate his ideas. On the other hand, a knowledge acquirer is better prepared to help when questions are more specific; specific questions imply that the knowledge acquirer has some corporate knowledge, both in terms of being able to construct the question and of being able to comprehend the answers. The requisite grounding can be obtained via unstructured interviews, but they can also produce reams of transcripts that may be difficult to follow.

A more structured method for acquiring knowledge is through the use of properly constructed questionnaires. If necessary, the expert can take time to reflect on the questions. Questionnaire questions are liable to the same shortcomings as those used in interviews, and responses can be even more superficial, but questionnaires do confine the expert's response, and they give a knowledge acquirer some basis for constructing more probing questions.

Determining whether to use a questionnaire depends on the topic as well as the depth of inquiry that is desired. Like interview questions, questionnaire questions can be directed toward eliciting general or specific kinds of information. Unlike interviewing, they are not a natural exploratory device: questions are pre-set, which assumes a topic has been settled on. However, in another sense, questionnaires can help orient the knowledge acquirer, by providing fundamental background information. It is a safe assumption that all occupations have certain abstract and concrete characteristics, such as work goals and work materials. Although questionnaires may not thoroughly plumb these topics, they raise questions that are easily answered, and give the knowledge acquirer a place to start.

Language should be taken into careful account when a questionnaire is used. The knowledge acquirer may not be available to clear up misunderstandings. Linguistic conventions and social rules vary from group to group, and concepts are named differently. For example, "lower class individuals 'see' but they don't 'perceive'" (Kinsey et al., 1965), and naval aviators "prioritize," rather than "order." Whenever possible, an expert should be involved in constructing the questionnaire. At the least, it should be reviewed by one before it is distributed.

Whenever possible, questionnaires should be distributed in person. This personalizes the exercise and gives the knowledge acquirer a chance to explain why they are being given. Experts may wonder why such simple information is being sought. Hence, it is important to explain the underlying rationale for using them, i.e., that they are useful for acquiring basic information but do not take much time to complete. Questionnaires need not be completed on the spot; experts may want time to think about the answers, which is the advantage of a questionnaire. Questionnaires should always be reviewed and pretested against criteria of repetitiveness, specificity, clarity, directness and simplicity. A poorly designed questionnaire can jeopardize goodwill and project credibility.

A potential problem with questionnaires concerns biased responses. Questionnaires, like interviews, can crystallize beliefs and opinions. Questions can be instrumental in changing vague and loosely held thoughts into cohesive and meaningful patterns of response, by prompting an expert to think about an issue in a new way. Such a crystallization of thinking is not altogether undesirable, but the results may be biased. They should be validated by other methods.

Open-ended Questionnaire. To get a sense of their value, two types of open-ended questionnaires were used by the ASPA knowledge acquisition effort. One questionnaire was directed at a high-level conceptual overview of the problem. Sample questions were:

- What are the major objectives in an airstrike plan?
- Why are these important?
- What constraints could prevent you from achieving your goal?

- What is the order in which you plan things? Why?
- What part of the planning is most critical? Why?

These questions evoked a very high-level frame of reference in the responses, e.g., to make political statements, to act as ambassadors, etc. The answers took on essay proportions. While the responses were informative, the expert, who was highly conscientious, was advised to stop after an hour and a half of writing.

Other open-ended questionnaires were designed for short answer responses. They were intended to elicit descriptive information, rather than explanations. The idea was to then incorporate the descriptive information in questions for use with more probing methods, such as interviews or walk throughs. The goal was to build a context for experts, which would help them answer "how" and "why" questions later.

These kinds of questionnaires were used in relation to a couple of topics. Their specific wording was tooled by experts. The types of questions included:

- The source of expertise, i.e., training, experience, schools?
- The first thing considered when doing an element of the planning?
- The most important information required to do that element of the planning?
- The identification of the information's source?
- The names of manuals and references that are most helpful, essential?
- The names of decision aids that are used?
- The persons (roles) consulted?
- The most complicated situation for that aspect of the planning?
- The worst and most common error associated with the task (critical incidents)?
- The least concern when doing that element of the planning?

Judged by the time it took the experts, the questions were easily answered. The fact that some experts' opinions differed was not a concern. The object was to obtain the fundamental information necessary to ask more specific questions.

Forced-Answer Questionnaires. A forced-answer questionnaire specifies both questions and sets of responses among which the respondent can choose. It assumes sufficient knowledge of a topic to permit construction of the questions; hypotheses about a topic are operationalized

by transforming them into statements with which a respondent can agree or disagree. Forced-answer questionnaires can yield more accurate results than open-ended questionnaires when they call for a respondent's recognition of a true statement rather than his recall. Consequently, they are a useful way of validating knowledge, although not as useful for initially gathering knowledge.

The ASPA project used a mixed version of a forced answer and open-ended questionnaire to validate data. What was notable about this questionnaire, discussed above under interviews, was its response format. Responses were restricted to "always," "never" or "sometimes," with "sometimes" requiring an explanation. This response format violated specific conventions of questionnaire design. Respondents normally resist being forced into absolute categories, and it is a general rule that sufficient and specific gradations of intensity (or frequency) should be provided as possible answers. For instance, "five times a month" is preferred to "frequently." However, the object of the questionnaire was to provoke respondents to search their long-term memory for exceptions, and report them. Violating design conventions served that purpose. It was important, however, to explain to experts that our intent was not to antagonize them, but to ensure that statements of "fact," which would be represented in ASPA's computer database, were indeed facts. Samples of true/false questions were:

- Weapon types are mixed within a single rack.
- Load inboard stations first to reduce drag.
- Minimize the number of racks used.
- Missiles or rockets can be loaded on station 3.
- Station 3 should be loaded last.
- MERs are preferred to TERs because they are more available.

WALK-THROUGHS

A walk-through is a method that entails having an expert "walk" a knowledge acquirer through his activities as he accomplishes a task. Walk-throughs can take two forms. The activities can be observed and noted, as the expert explains what he is doing. Alternatively, the expert can instruct the knowledge acquirer, telling him each step to take, as if he were teaching him to do the task.

A walk-through is conducted in the expert's own setting, where work materials and the environment provide memory cues. The task itself focuses the knowledge acquisition process. A walk-through does not require an expert to think in a vacuum and additional memory aids, like files, are easily accessible. Because it does not pull an expert away from the job, it is a relatively unobtrusive method.

This method was attempted a number of times, with varying degrees of success. Different experts walked members of the ASPA team through a weaponeering problem, using various decision aids and manuals. The walk-throughs succeeded in making the point that there were different technical approaches to the problem. The main products of the method were identification of manuals and decision aids, familiarization with their uses, and obtaining sample outputs. The samples were important later as the basis for read-throughs.

Walk-throughs can be difficult to conduct and, as with other methods, important points can be missed. The speed with which experts report can be a problem. The appropriate candidates for a walk-through are those most accomplished at a task, or, as in our case, the only ones capable of doing the task at all. For them, doing the task is second nature. They can go through the steps very quickly, and the highly technical language used in their description make recording the information difficult. Tape recording is not the entire solution because the pertinent information may be visual, or may involve physical activities such as entering data into a computer.

A walk-through may not capture a task's links to a larger problem, either conceptually or with respect to information flow and use. Even though walk-throughs take place in natural settings, the walk-through itself is an artificial situation. A boundary is placed on the example task in order to focus on step-by-step procedures and, consequently, the procedure's relationship to a larger process may not be entirely evident. For example, by focusing attention on weaponeering, weaponeering's larger relationship to the process of airstrike planning, which involves tactics, aircraft capabilities, etc., may not come through. That kind of information may be so obvious to the expert that he may not mention it during a walk-through, especially when he is focused on the task at hand. Furthermore, if the knowledge acquirer lacks a prior understanding of the conceptual links (the flow and meaning of information), and is fully employed in following a complex task, it may not occur to the knowledge acquirer to ask specific questions about those issues.

Walk-throughs are also artificial insofar as a real problem is not being solved. The expert may use imaginary information that would have been received from elsewhere in a real planning scenario, or may apply special expert knowledge; it is not always clear. A walk-through does not cover all of the situations that would actually occur. That would call on different knowledge and problem solving approaches. Finally, because an expert is a specialist at a task, the difficulties that the less accomplished may have may not come through. For instance, until the information was specifically solicited, the ASPA team did not find out what the vast majority of aviators disliked about their decision aids and why they avoided using them. That information pointed to the mistakes that should be avoided in a new system.

The checklist given below is based on the experience of the ASPA team. It might be helpful to have the expert review the checklist before a walk-through, to enhance understanding of the kind of information that is desired.

- Define the situation for the expert. The object of a walk-through is not to have a demonstration of the expert's job, but rather to set up a situation where the expert can verbally describe what is being done while doing the task, so details can be recorded and clarifying questions can be asked later, when they are not disruptive.
- Identify the objectives of the task.
- Identify the source of information brought to the task. What information and priorities are brought to the task? Does it vary? Where does it come from? Where does it apply in the procedure?
- Find out if the expert's way of doing the task is typical or a short cut.
- Obtain the material product of the procedure. Whenever possible, material samples of the product of the procedure should be collected. They make more sense later, as knowledge accumulates. They are also a good material grounding for asking questions.
- Identify constraints on the procedure. What circumstances would alter the way in which the task is done? This is a very general question; an example or two can probably be obtained.

READ-THROUGHS

"Read-through" is a name for a method originated at the Naval Ocean Systems Center. The method entails having an expert instruct a knowledge acquirer in how to read and interpret documents. Read-throughs were inspired by the fact that military documents tend to be massive, numerous, obscurely written, and inconsistently formatted. The experts themselves are sent to school to learn to read them, and they may be distinguished as specialists largely on the basis of their mastery of manuals. In fact, the problems associated with the use of these manuals made the prospect of an expert system, which could integrate the information in a more accessible way, highly attractive to operational personnel.

A read-through entails having an expert describe the actual steps taken when using the documents. An expert defines the kind of information that is used, where it is located, and the order in which it is used. Directions or checklists included in the manuals themselves are often ignored in reality. Directions for using series of manuals are not available. Mental short-cuts, identification of (some) obsolete data, and a sense of the context in which the information is used come through in read-throughs. The documents are usually familiar to the expert and provide memory cues. The situation is focused for the knowledge acquirer by the documents. Less demanding on both parties than walk-throughs, read-throughs can comfortably be quite lengthy. Sufficient time and access to appropriate documents are the basis for successful read-throughs.

Hard copy readouts from decision aids are another candidate for read-throughs. With these in hand, it is easier for an expert to identify and define his inputs, and identify where they

come from. It is also easier to identify and define the outputs, and to elicit an expert's explanation of what they are related to and used for.

Because manuals provide effective memory cues to an expert, they can be especially useful when the knowledge acquirer knows enough to have a sure topic and line of questioning in mind. A list of objects, for instance, helps an expert remember the objects' associated characteristics more easily than if both the list and the characteristics had to be generated. For example, a list of weapons, somewhat obsolete, enabled an expert to quickly name the tactics that could, and could not, be used with the individual weapons.

The problems with read-throughs mainly involve locating the appropriate documents and course materials. There are so many manuals that even identifying them was a major knowledge acquisition undertaking. Obtaining them took considerable time, and some were outdated. Another problem, which reflected on the manuals, not the method, is that so much of the data in even the latest versions of manuals is considered suspect or inappropriate by operational personnel. Making the best use of manuals as direct data sources requires a good deal of expert time, and access to the experts can be difficult.

OBSERVATION

Observation is a field study method that is often used to explore a setting, such as an organization, before enough is known to settle on a specific research problem and to construct a hypothesis (Dalton, 1967). Observation is also used to gather empirical data to confirm working hypotheses derived in laboratory experiments or from tests, questionnaires, or interviews, since actual behavior often differs from the results obtained by those methods. Observers often notice things that the people working in an environment take for granted or are unaware of (Brandt, 1972; Brayfield and Crockett, 1955; Cicourel, 1964; Wernimont and Campbell, 1968; and Wicker, 1969). Observation is also used to enhance the investigator's understanding of the cultural context associated with different points of view (Blau, 1963; Dalton, 1959; Gouldner, 1954; Selznick, 1949). Observation is less obtrusive than a walk-through. The observer watches behavior or activities as they proceed naturally, and then asks questions.

Another important difference between walk-throughs and observation is the focus of the knowledge acquirer's attention. Observation can be focused on the activities of an individual or of an entire group; a walk-through requires focus on one individual, who "talks through" a task.

Like interviews or questionnaires, the method can be highly structured or unstructured, depending on the degree to which the topic of investigation has been decided, though here structure resides in the decisions regarding what to look for, instead of what to ask.

The ASPA team's attempt at observation was exploratory. The observation was structured loosely by five goals:

1. To confirm our working hypotheses regarding airstrike planning, based on interviews with experts.

2. To identify specialists;
3. To identify the problems those specialists dealt with;
4. To identify the sources and flow of information; and
5. To identify characteristic reasoning associated with airstrike planning.

The observation was conducted on practice planning sessions in an academic setting. The practice sessions were almost facsimiles of real planning sessions, except that more time was spent discussing the background and purpose of each step.

Direct observation of the development of an airstrike plan is not easy. Quite a number of people are involved. Experts are briefed, assignments are worked out based on an initial assessment of the problem, and the different working groups disperse. Fortunately, faculty were available to direct us to important areas, answer questions we raised, and explain what was typical of airstrike planning and what was an artifact of the academic setting. We were able to see what kind of information planners typically receive and we could look at the physical products of their planning efforts. Examples of typical reasoning could be overheard in the planners' discussions among themselves. Observing the briefing sessions preceding and following the planning sessions, and listening to the accompanying debates, contributed to our basic understanding of the factors and reasoning associated with real-world scenarios. The plans developed during these training sessions were later evaluated by both faculty and electronic means.

Observation corrected misleading impressions that we had observed from the experts' verbal, and necessarily idiosyncratic, descriptions. It provided an overview of airstrike planning and gave us a solid framework for constructing questions. Observation also confirmed our working hypotheses about airstrike planning, i.e., that the problem to be solved and the availability of information will affect the elements that are included in a plan, the priority given to problems, and the expertise called on to address the element. Further, specialists proceed fairly independently with problem solving, and then reconvene to resolve incompatible features of their plans. Finally, being on site for a time gave us the opportunity to collect checklists, course materials, the experts' own flow charts and other documents, and establish important relationships.

Like the other methods that have been discussed, observation is not suited to achieving all knowledge acquisition goals, and there are practical limitations to the method. Access to the people and places to be observed is foremost among the problems. It took the ASPA knowledge acquisition team a considerable length of time and involved considerable effort to gain access to a suitable site. The degree to which observation is unstructured, and the degree to which the knowledge acquirer must remain responsive to the environment, may make inexperienced knowledge acquirers uncomfortable, even more so than during unstructured interviews. It is not a method designed to plumb issues deeply, but rather to let the situation speak for itself.

In order to overcome the physical problems associated with observation (i.e., the mobility of planners), the ASPA team obtained agreement from the staff to conduct a one-room planning session, where the planners would stay in one place. Some of the natural features of the situation were distorted, but we were already aware of those features. The session was recorded on tape while it was being observed by the knowledge acquirers. Group interviews were conducted later and specialists were interviewed individually to clarify questions generated by the one-room session.

RECORDING THE DATA

Just as an expert cannot be relied on to remember what is known without memory cues, neither can a knowledge acquirer. Recording the data acquired from experts is essential. Tape recording is the method of choice, because so much can be missed in notetaking alone, especially when the knowledge acquirer is not fully conversant with what is being said. But sometimes there is no choice but to take notes, for instance, during a walk-through where much of the information is visual. Also, some people become guarded and tense when confronted with a tape recorder or video machine. An expert's permission should always be requested prior to being recorded but even when received the knowledge acquirer should remain alert to the possibility that taping might have a negative effect on the interview.

Thorough notetaking slows down the interview and can be disruptive to some expert's concentration. However, depending on the characteristics of the expert and established rapport, slowing the interview to a less than normal conversational pace is not always a problem. Pauses allow time for reflection and experts can be quite considerate of the fact that a conscientious effort is being made to record what they have to say accurately. Reading notes back to them underlines the point and gives them the opportunity to straighten out inaccuracies: what was "heard" and recorded is not always what was said and, upon reflection, the experts may find that what they said was not what they really meant.

DOCUMENTARY SOURCES

Technical Documents. Many technical and academic documents obtained for the ASPA were not very helpful in themselves, prompting "read-throughs" with experts. However, the effort devoted to understanding them sufficiently to produce and solve sample problems and to produce task descriptions was fruitful. Some technical vocabulary was mastered and concepts, such as "fragmentation envelope," were effectively introduced with photographs and charts. The complexity of the manuals, and the time required to solve problems with them, was demonstrated. Later, expert comments such as "I would like to see a computer do the coolie work" made sense.

Magazine Articles. The experts' own magazines and articles were much more fluent than the manuals, and effectively communicated the fundamentals of airstrike planning and the "gouge," or the latest rules of thumb honored in the expert community. The language and

concepts were specialized, but the context provided in the documents, and the fact that knowledge gaps could be directly pointed to, made these kinds of documents much easier to decipher.

Training Courses. Attendance at some courses, video tapes of others, and course notebooks also effectively introduced concepts and provided some graphics, although some courses were inaccessible. Participation in the courses produced more conceptual understanding than rules. It provided background information essential to pondering course materials and other documents, which introduce specific objects and their characteristics. These kinds of documentary sources were especially helpful because the codified knowledge was presented in a format specifically designed to teach and discuss. The major problem is that too few such documents are available.

Experts' Flowcharts. Two of the airstrike community's own attempts at modeling the airstrike problem were obtained, and are presented as figures 1 and 2. These flow diagrams capture different dimensions of the problem: the division of labor associated with tasking, with intersecting decision points, and conceptual variables that bear on decision points.

FORMALIZATION OF KNOWLEDGE

Classification of incoming data for the ASPA project was unexpectedly difficult, reflecting the magnitude and complexity of the airstrike planning domain. The results of preliminary interviews and document searches all have ongoing relevance to knowledge acquisition, interface development, and hardware and software selection. The data's classification, or coding, affects whether it can be accessed for those purposes at appropriate times, or is functionally lost. Especially when a team is working on the project, standardizing categories for filing data is fundamental to making the knowledge manageable and useful. How the information is stored, whether in paper or computer files, affects its accessibility to various team members. Unless the data is documented and mutually accessible, different team members will "know" different things and when team members leave the project, corporate knowledge will be lost.

A number of methods for formalizing the data have been tried by the ASPA knowledge-acquisition team. Early in the project, in order to prototype and test system architecture, the small subdomain of weapons loading was used as the focus of a trial knowledge acquisition. Functional diagrams, task descriptions, and task-flow diagrams were used to formalize the knowledge. These three methods have been described in detail in an earlier report (Ehrler et al., 1985). Briefly, a functional diagram describes the high level conceptual factors which bear on the problem. A task description focuses on verbs and step-by-step procedural aspects of a problem; it describes what is done and can vary in its level of detail. Task-flow diagrams graphically display the decision-making flow, specifying what shall be left to the user, and what given over to the computer. The values and shortcomings of these ways of formalizing knowledge have been reported earlier: for instance, that a task description is difficult for an expert to comprehend and evaluate. Besides those limitations, the latter two kinds of formalizations can be somewhat static and linear; what is done is conveyed, not how or why. Consequently, the computer-science team members, who were looking for the basis of rules in task descriptions and flow diagrams (explanations) were unimpressed.

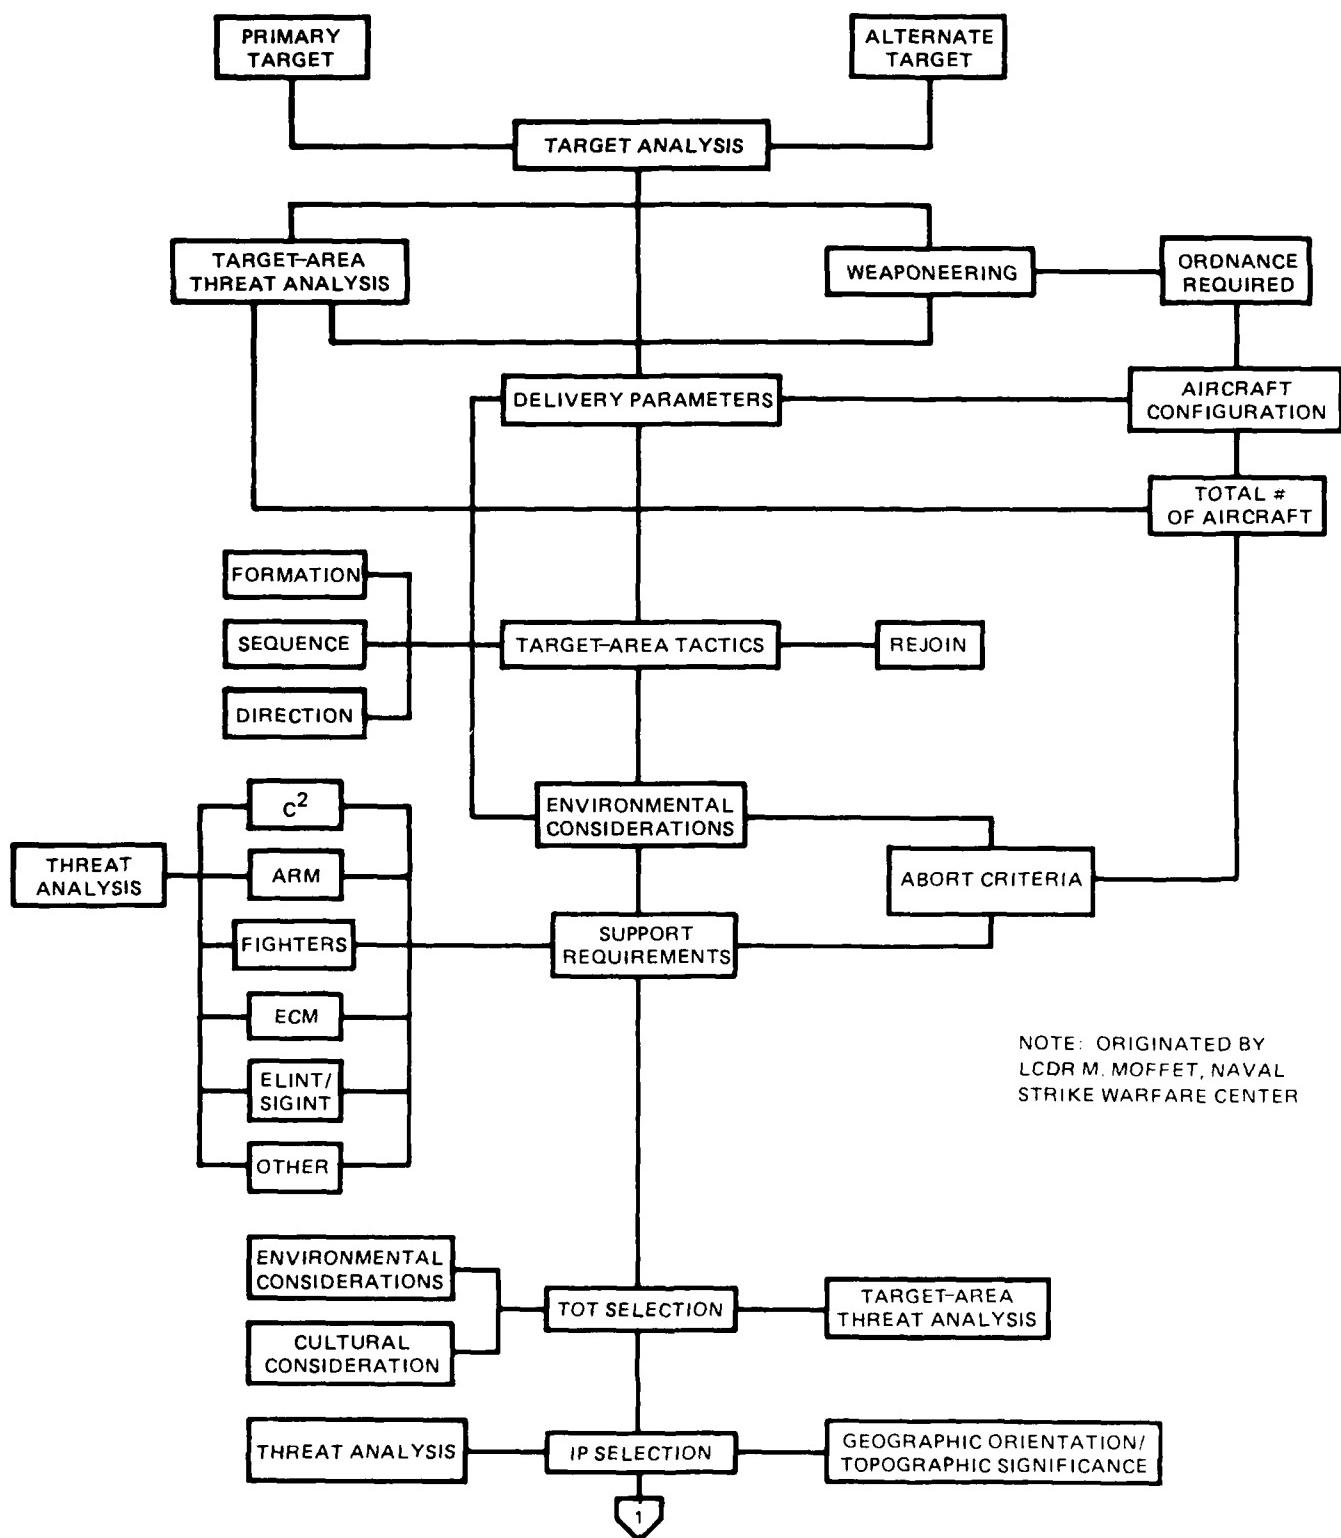


Figure 2. Airstrike planning sequence (page 1 of 2).

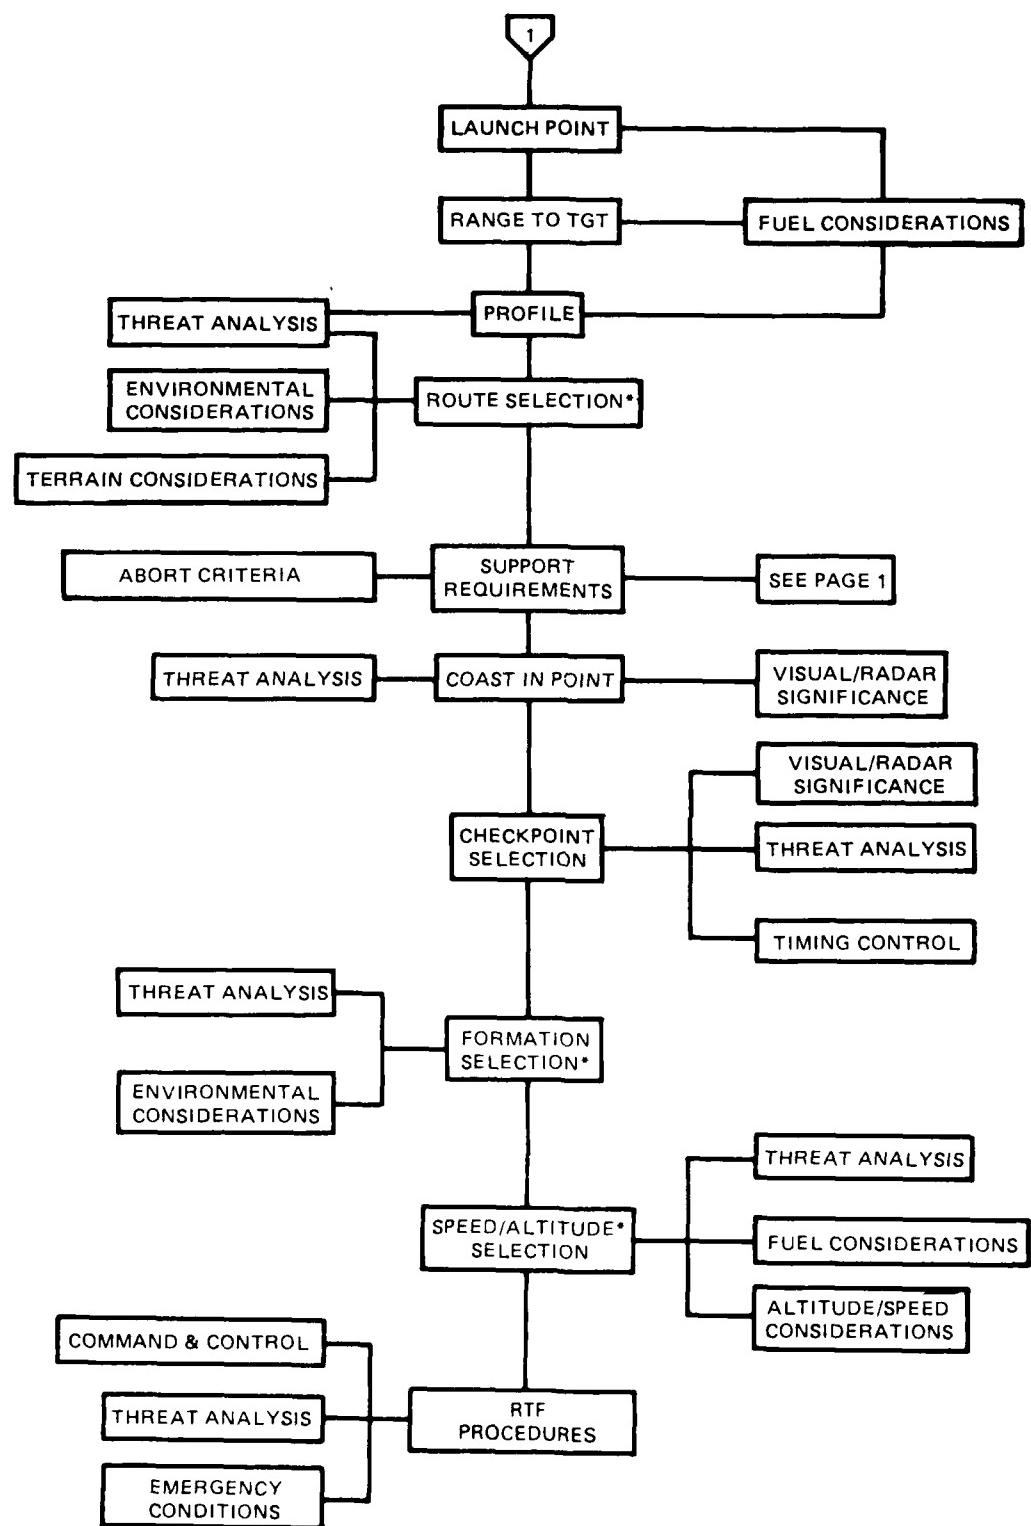
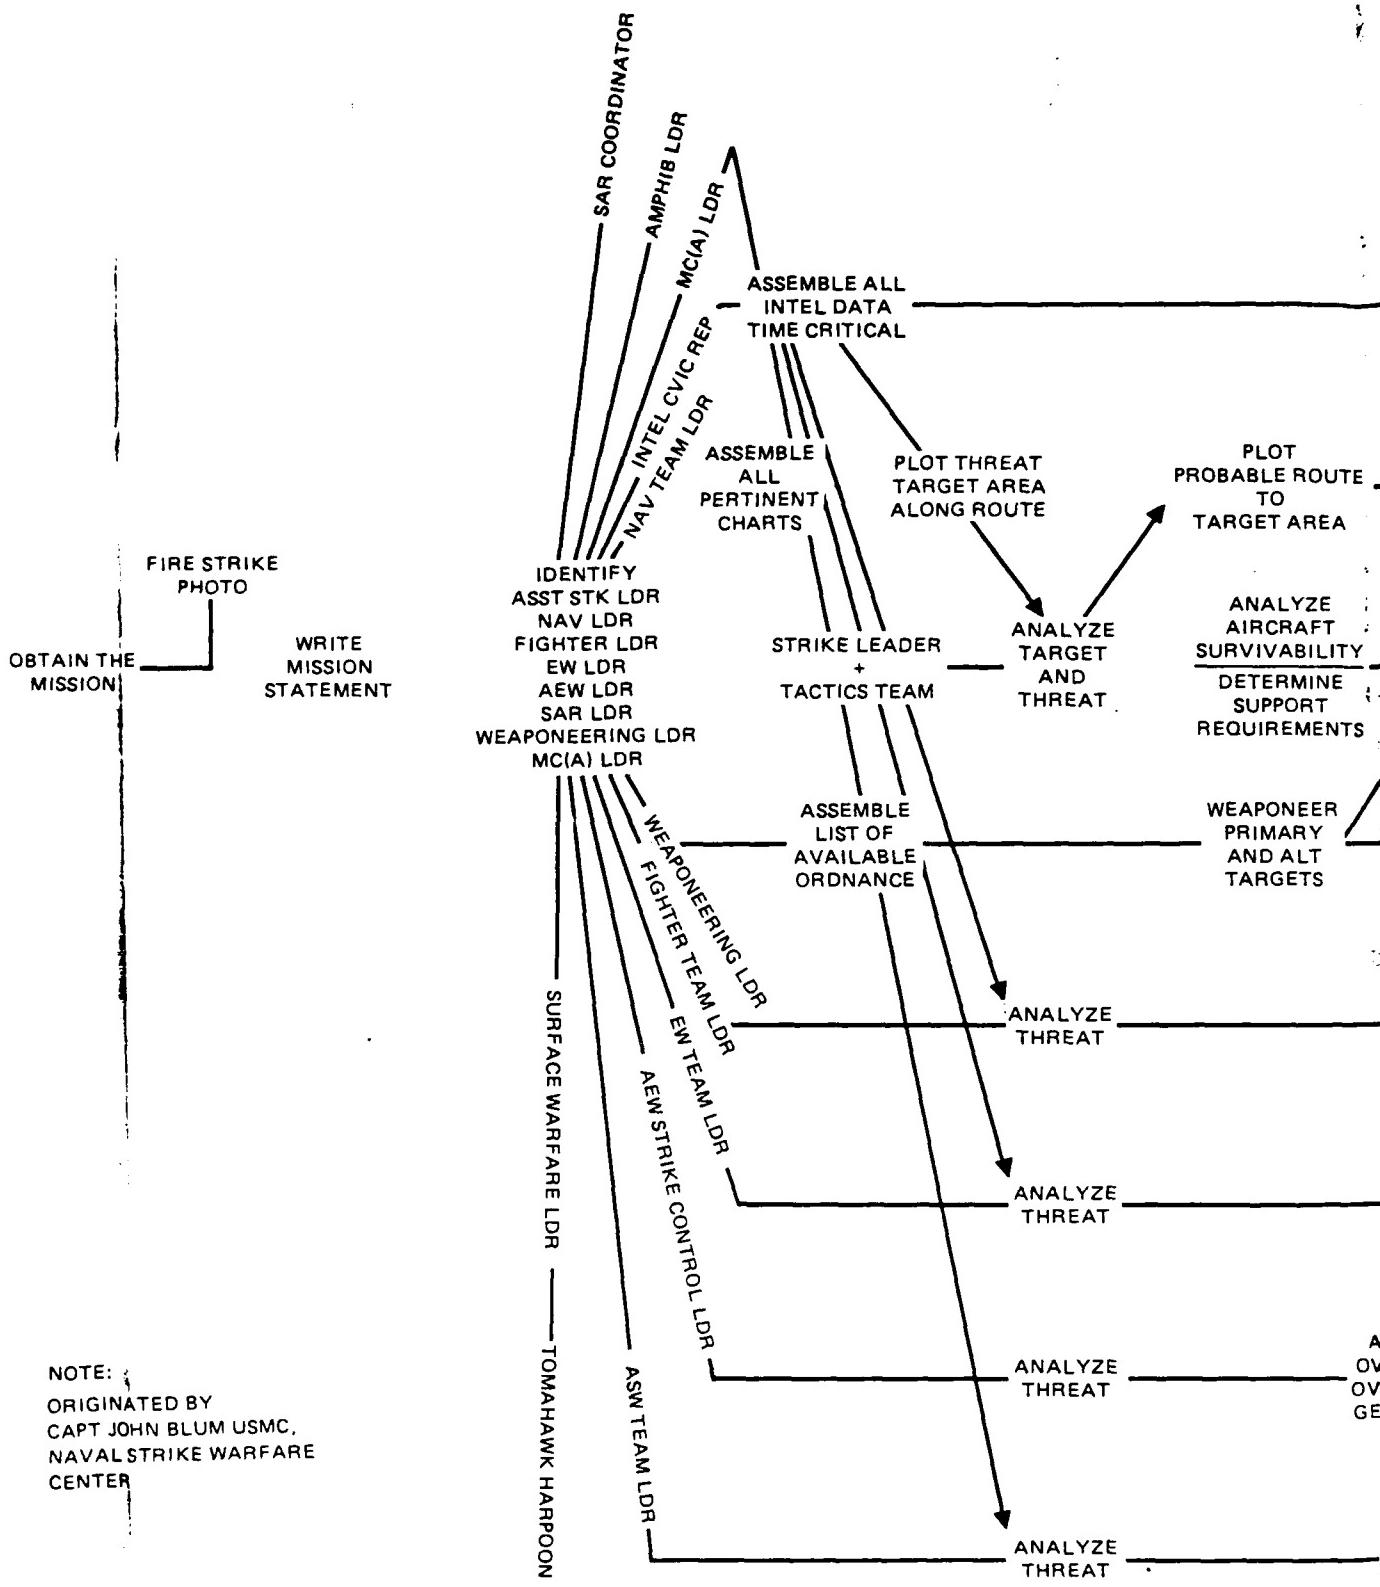
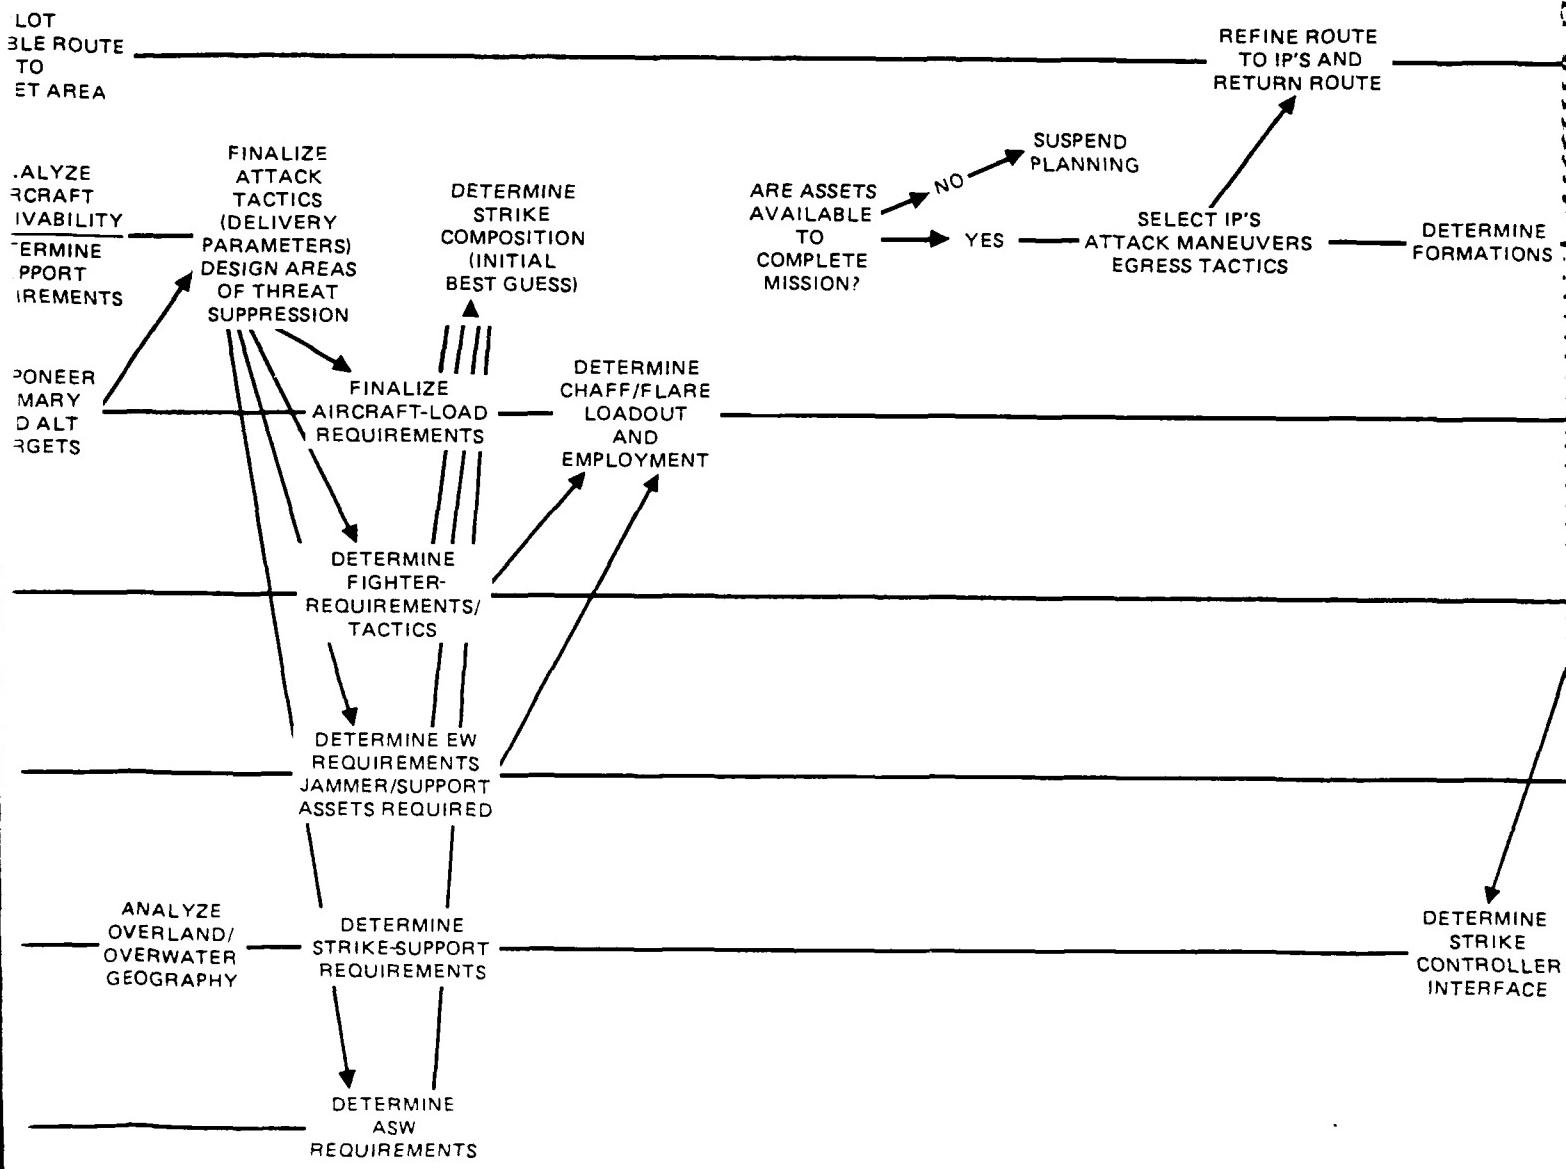
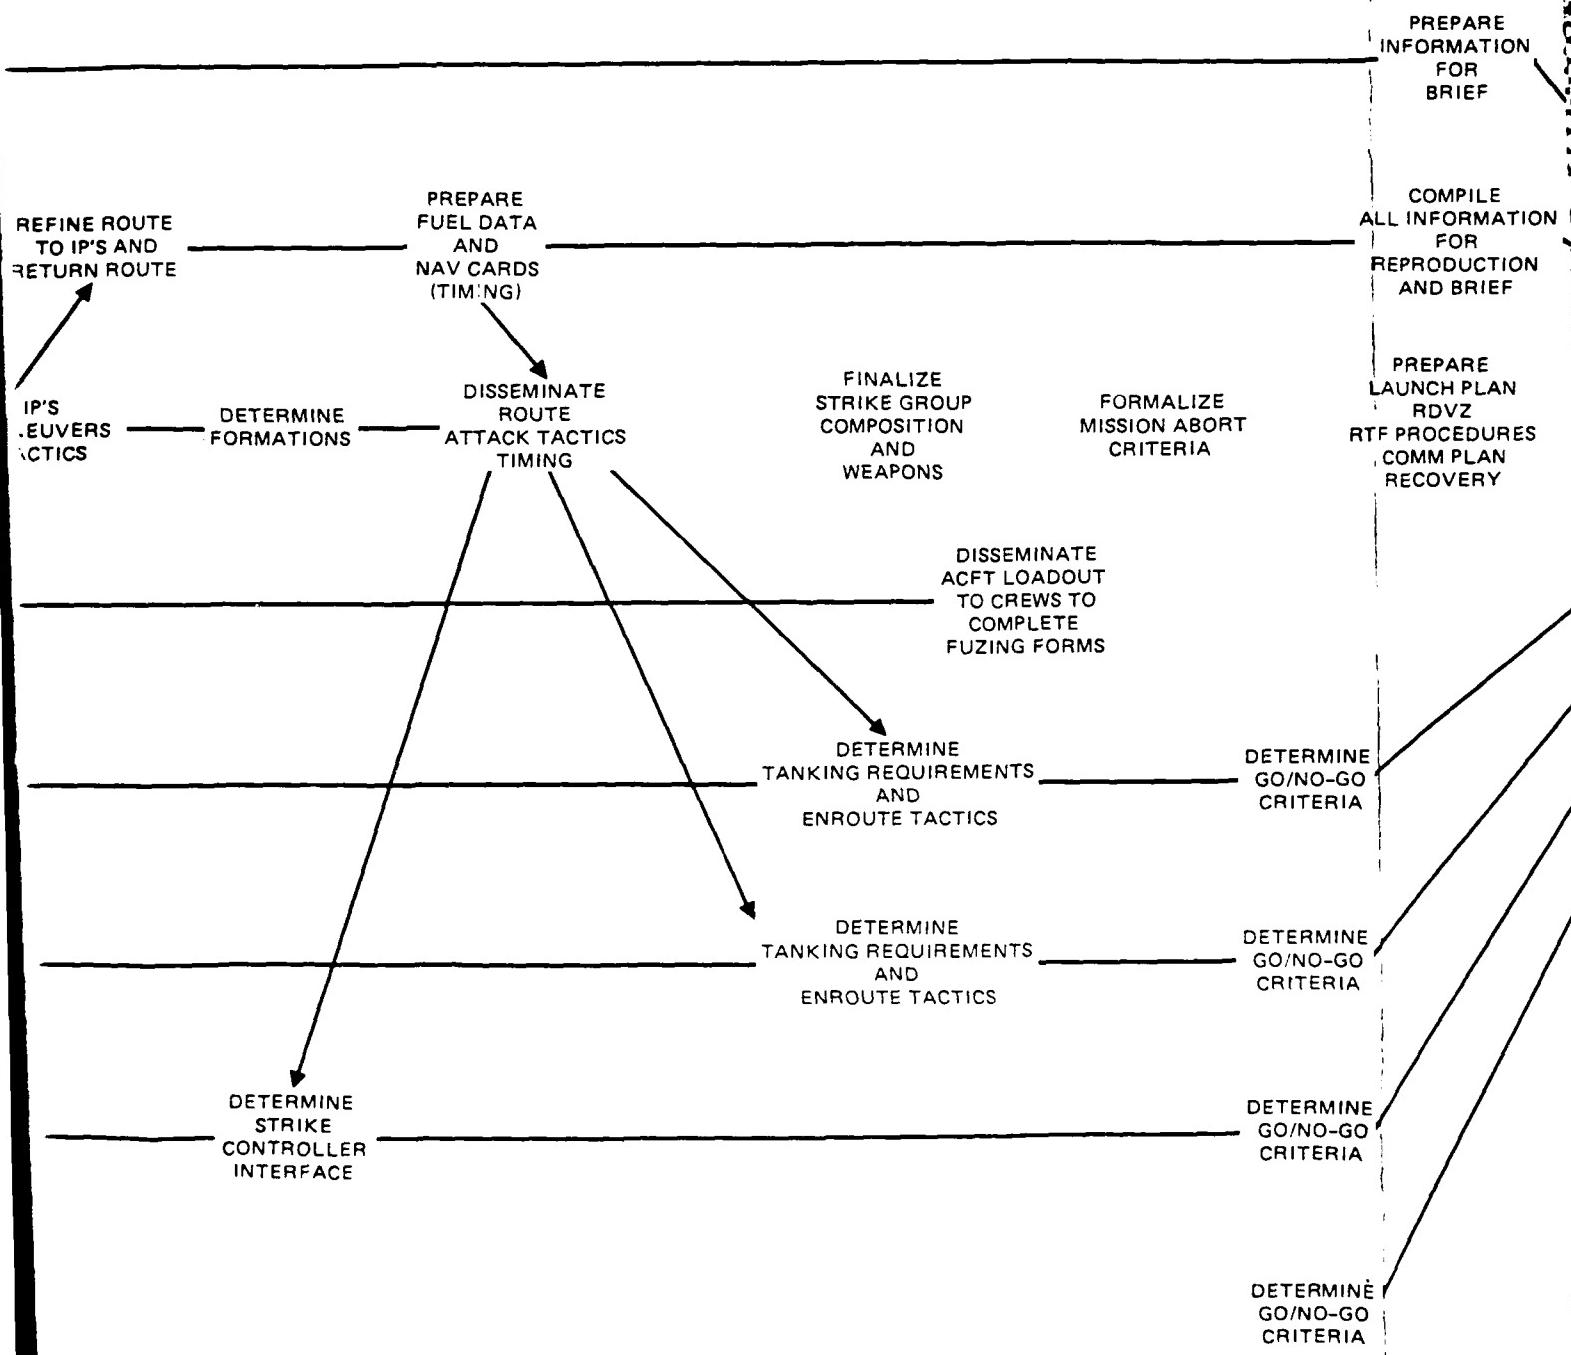


Figure 2. Airstrike planning sequence (page 2 of 2).







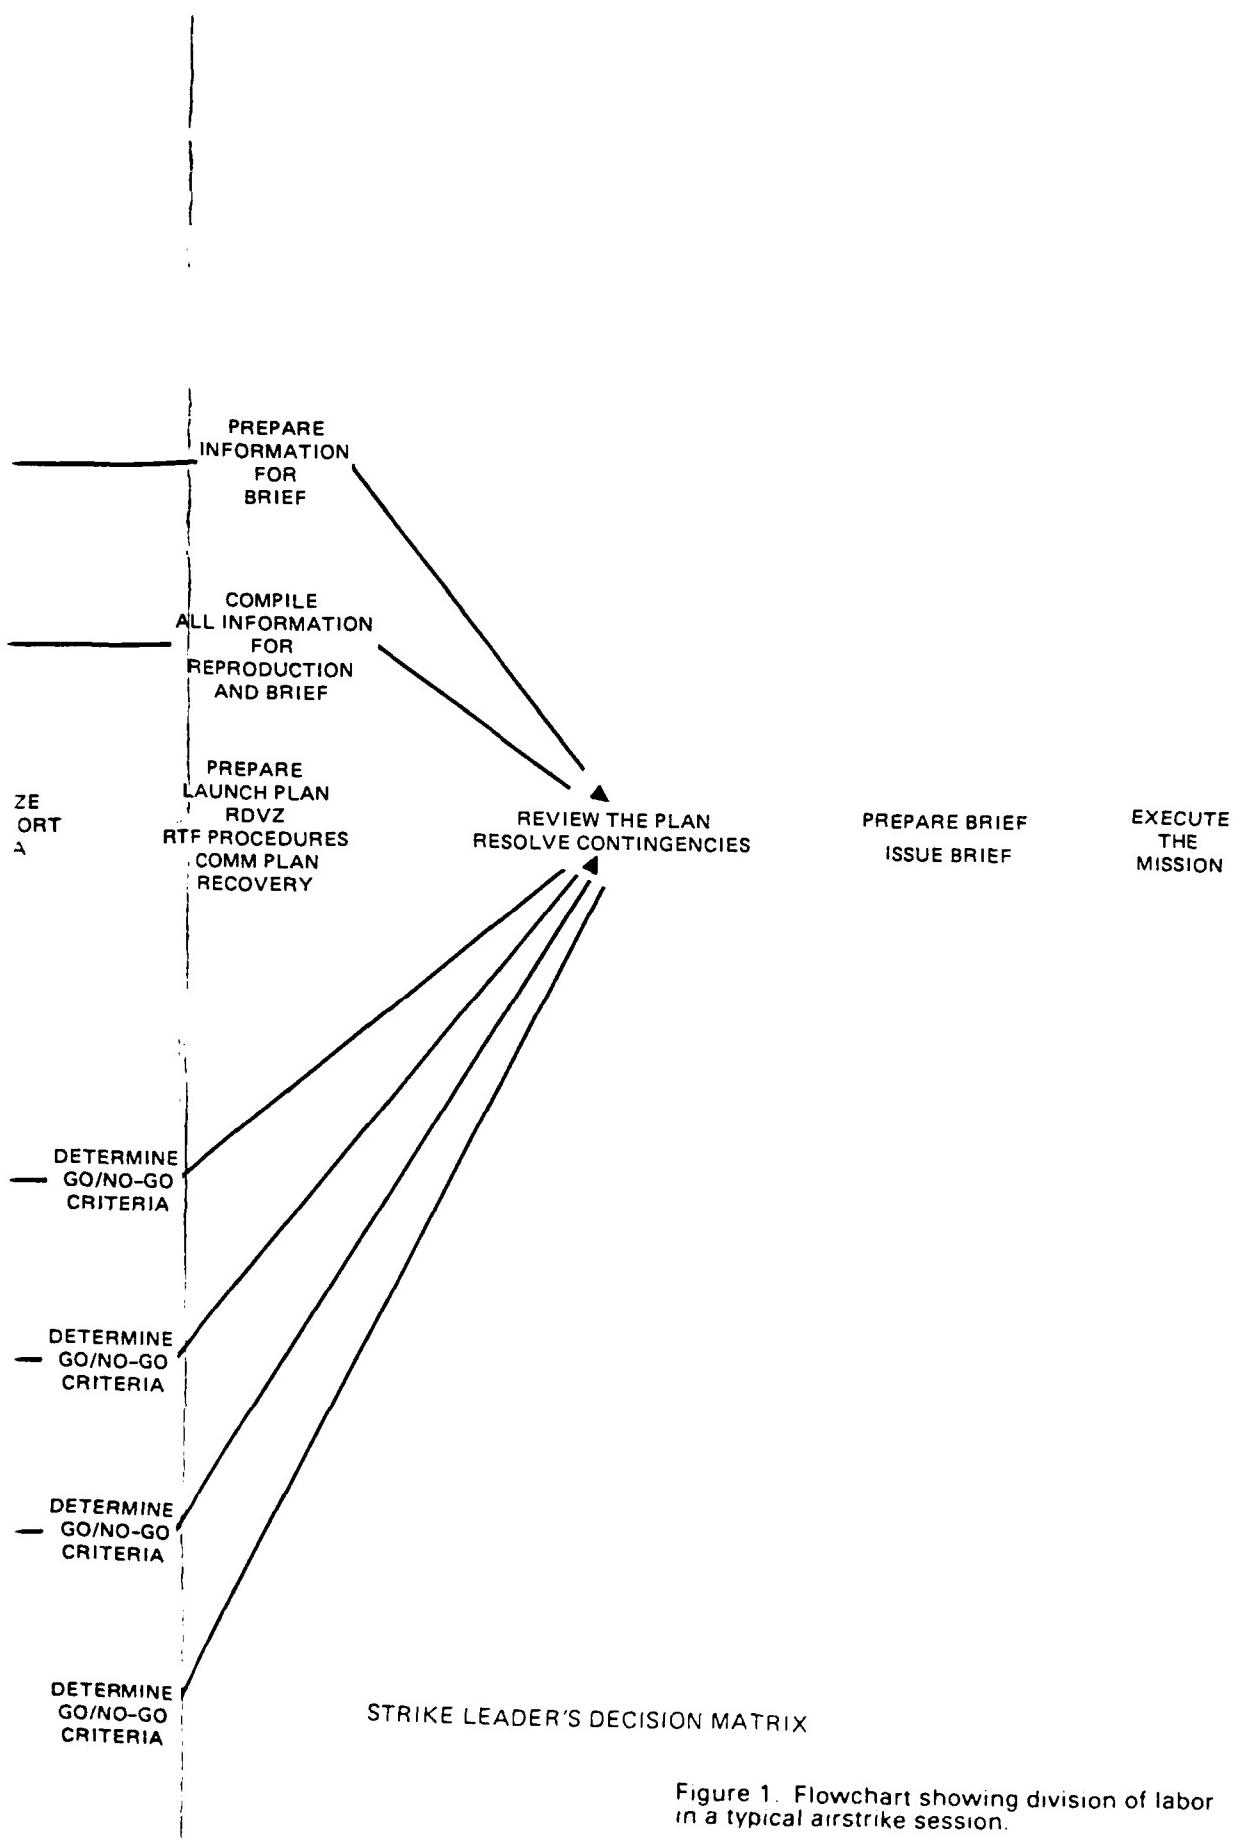


Figure 1. Flowchart showing division of labor in a typical airstrike session.

When the knowledge-acquisition focus moved away from a specific topic, weapons loading, to a new subdomain, different problems emerged. In order to select a new subdomain, airstrike planning itself had to be understood well enough to support the selection. Another knowledge acquisition strategy, using a simple scenario as a focus for acquiring knowledge, was also envisioned. The exploratory-knowledge acquisition demanded by those goals produced a deluge of information before a model was conceivable, and even before categories for filing the information became clear.

A compendium of knowledge was created, consisting of a loose-leaf notebook broken into broad headings, such as weaponeering, target-area tactics and route planning. This compendium, or notebook was inspired by a number of facts. Access to experts is at a premium in a military domain and all of the information that is gathered is valuable, even if it does not exactly pertain to an immediate knowledge acquisition goal. A great deal of information and data are required to understand the domain, such as vocabulary, regardless of whether or not it is directly represented in the software. Accumulated knowledge is the basis for framing probing questions that lead to programmable knowledge. However, not all of the required knowledge can be acquired, or programmed, at once, but needs to be saved and reviewed. Compiling the information in a notebook gives experts the opportunity to correct and elaborate the accumulated knowledge. It also gives new project team members the advantage of seeing the corporate knowledge base.

Documents and notes were culled for the conceptual elements that affected planning for weaponeering, target-area tactics, and route planning. The elements were then laid out in tables, which were labeled "input," "output," "constraints," and "rules." An expert then put the elements in priority order, elaborated some elements, moved some elements from one category to another, and even suggested fruitful lines of questioning related to areas that were not his speciality. Those tables became the "table of contents" for our original topics. Notes, documents, interviews, flow charts, etc., which were related to the topics, were filed in the compendium. New topics were added as they emerged. A glossary and typical reasoning scenario were also included. Although not exhaustive (manual and other documentary sources were not included), the compendium functioned somewhat as a resident expert, against which the programming effort could check its direction.

KNOWLEDGE ACQUISITION CHECKLISTS

The problems with knowledge acquisition have not been solved by the ASPA project, but many have been clarified. Many difficulties were connected to the fact that collecting knowledge in a military setting, for a military application, is very different than producing a traditional system. However, guidelines that addressed those differences were not available. Because producing a system for military application can be quite involved, given the complexity of the domains, some suggestions for how to begin are offered here.

The first problem is to get oriented in many senses. Finding the appropriate expert candidates ranks high on the list of initial tasks. The following general questions should be asked:

- Who knows about the domain?
- Who are the reputed authorities in the domain?
- Where are the most accomplished authorities physically located?
- What kind of organizational setting are they located in?
- What are the values in the setting?
- Who makes formal and informal decisions in that setting?
- Can time be devoted to a project? On what basis?

Besides finding the appropriate experts, the next impediment to knowledge acquisition in a military domain will be encounters with new words, objects, acronyms, symbols, and concepts. Because exhausting an expert and jeopardizing project credibility is ill-advised, and textbooks may not be available, the following strategies for overcoming the initial communication hurdles are suggested:

- Talk with those familiar with the field.
- Tour the material facilities.
- Pictures are worth a thousand words--browse manuals for them.
- Read professional magazines.
- Read related newspaper articles and novels.
- Watch television documentaries and movies.

- Obtain samples of what is produced.
- Obtain course curriculums.
- Obtain job descriptions.

Although a full consensus exists in the literature that the first goal of knowledge acquisition should be to understand the domain problem, this is not a straightforward undertaking, and is even less so in a military domain, where there are problems that may vary considerably in their individual structure. The following strategies are suggested as guides for locating the problem(s):

- Discrete, high-level elements of the problem should be identified. What areas do experts talk about that seem fairly independent?
- What is the division of labor? Is it static, or does it emerge in response to factors? What factors?
- If it is a planning problem, does the planning take place sequentially or independently? Is it a linear problem or a set of problems?
- What is it necessary to know to proceed with an element of the planning/problem solving process.
- Is there a logic of priorities (e.g., survivability) that links or orders problem elements or alters tentative solutions?
- What parts of the problem do experts believe are most amenable to a computer solution (they may not be familiar with expert systems).
- What parts of the problem do experts find most tedious and time consuming?

DISCUSSION

Experience with the ASPA Project highlighted a number of problems that need to be solved before knowledge acquisition can be done effectively and unobtrusively in Navy field settings. The major problem areas concern how best to acquire knowledge from experts, how to control the knowledge effectively, and how to translate it into a form which facilitates programming.

ACQUIRING KNOWLEDGE FROM THE EXPERT

Access to credible expertise was a roadblock early in the project. Once access was gained to sufficient numbers of experts in the various specializations, the emphasis shifted to assessing different methods for drawing out their knowledge. Two problems guided the assessment efforts. The first problem was to find out if, as relevant literature suggested, some knowledge acquisition techniques were more effective than others in helping experts express their knowledge, work processes, and reasoning.

The second problem was trying to find out if there were features of knowledge acquisition processes that should be used when selecting knowledge acquisition methods. Could any links be found between the use of different methods and the types of knowledge their use produced? If use of a method did influence the kind of knowledge obtained, did the knowledge correspond to any analytically identifiable, necessary steps in the process of knowledge acquisition? Were some techniques more successful than others in producing knowledge that was particularly useful to programming efforts?

CONTROLLING THE KNOWLEDGE

Controlling the knowledge was unexpectedly difficult. The unstructured methods that were necessarily used in the first stages of knowledge acquisition as the domain was explored produced large amounts of data before members of the team were cognitively prepared to process it, or to record it in ways that were mutually comprehensible. Because of the complexity of the airstrike planning domain, a vast amount of knowledge was accumulated before a pattern appeared. When a subdomain was selected for testing the system's architecture, a substantial amount of general airstrike planning knowledge was accumulated in conjunction with knowledge acquisition for the subdomain problem. Even though the knowledge was not immediately relevant to the specific subdomain problem, weapons loading, it was relevant to an understanding of the overall domain problem, and warranted saving. But differences in the way various team members stored and indexed information made it inaccessible to other members. Finally, there was no definition of knowledge that satisfied team members; their intuitions and analyses reflected their differing project interests and disciplines.

In response to these practical issues, a major problem was finding ways to formalize knowledge so that it would be accessible to all project members when and where it was required:

for knowledge acquisition (to take into account corporate knowledge), for problem (or subproblem) modeling, for interface design, and for programming.

Since airstrike planning is a very complex domain, and since the ASPA project involved different kinds of work--and knowledge--there was a need for a lexicon for categorizing or describing knowledge. Software knowledge representation formats were considered too restrictive for three reasons. First, more knowledge was acquired than would be directly programmed. Second, determining the appropriate knowledge representation scheme is a function of the type of knowledge, thus formalizing the knowledge so it can be analyzed should come first. And third, obtaining an initial validation from an expert, to ensure that what is recorded is accurate, requires expressing the knowledge in a format intelligible to the expert.

A number of approaches have been taken to the formalization issue, for instance producing a loose leaf notebook compendium. A current approach, which has the most bearing on knowledge acquisition methods *per se*, is inductive. The idea is to base formalization categories on the kinds of knowledge categories that project experience has suggested are important, and link methods to those knowledge requirements.

Some preliminary steps toward this goal have been taken. A member of the human factors team at the Naval Postgraduate School is using three computer-based techniques for storing accumulated knowledge: a standard database management system, a text management system, and a graphic information representation system. This should enable the computer scientists to access accumulated knowledge when they are ready for it.

TRANSLATING THE KNOWLEDGE TO COMPUTER COMPATIBLE FORM

A major goal of the ASPA knowledge acquisition effort has been to find ways to expedite the translation of knowledge into computer software. The attempt at making knowledge acquisition systematic, by connecting methods to identifiable knowledge requirements, partially addresses this concern. Ideally, knowledge acquisition methodology should enable knowledge acquirers to turn over program-ready material to programmers.

The distillation and formalization of knowledge in ways that facilitate rapid programming involve more work. It would require that the human factors and computer-science teams become better acquainted with each others' tools. Having a better understanding of the knowledge representation schemes and problem-solving strategies that are currently available to artificial-intelligence work could enable knowledge acquirers to more rapidly analyze the programmable dimensions of knowledge, and translate it into that form. Computer scientists could profit by a better understanding of tools such as task analysis. This is because all of the knowledge that is required to properly understand a domain and its subdomains (which can differ widely in their structure) does not appear directly in software. For instance, flow diagrams provide a systematic way of doing function allocation, the identification of which problems should be solved by the computer and which should be left to the user. A task analysis provides a procedural overview of the problem at hand, a necessary step in understanding where more in-depth knowledge-acquisition techniques should be applied.

RECOMMENDATIONS

Experience from the ASPA project suggests that there are a number of kinds of knowledge that an AI project requires. Some approaches for meeting those requirements are suggested below.

OVERVIEW

Early in the knowledge-acquisition process, a simple overview of the problem which described elementary concepts and relationships would help orient the knowledge-acquisition team. If a problem definition can be constructed from available notes and documents, it should be verified by an expert; if an overview is not available, a questionnaire designed to elicit simplified statements of the problem should be constructed. An early goal in approaching a new domain or subdomain should be to acquire the kinds of elementary knowledge an introductory textbook would supply, if such a textbook were available.

VOCABULARY GLOSSARY

Understanding the expert's vocabulary proved crucial to the first stages of the knowledge-acquisition process, i.e., when interviewing experts. Experts use various names and acronyms for objects when giving explanatory examples. Definitions often are not requested during interviews, since such requests may disrupt the concentration of both the expert and the interviewer. Consequently, compiling a general glossary for the domain should be an early goal of knowledge acquisition to help interpret interviews. Detailed descriptions of relevant objects should be acquired later, when it is clear which objects will be included in the program.

PROCEDURES

Capturing expert procedures is the most difficult part of knowledge acquisition. To determine how a computer can replicate (or aid) expert problem solving, a faithful description of what the expert does is required. It is difficult to acquire such a description because knowledge, actions, and reasoning are all involved in expert procedures. Consequently, a combination of methods, designed to capture both the behavioral and cognitive aspects of the planning problem, is called for. The data derived from the use of the various methods should be interpreted, integrated, and formalized in ways that make the data cognitively accessible to both computer scientists and experts.

To identify and systematically capture the primary information categories with which an expert works, two methods should be used. Notes and documents should be reviewed for knowledge that can be used to construct a table of the classes of objects that constitute the relevant information base. This should then be presented to the experts for review and completion. A questionnaire should be constructed to enable the expert to identify what essential information he needs to accomplish the task, what would constitute perfect information, what constitutes typical information, and what is helpful, but not essential, information.

Another early goal, when working in a planning domain such as ASPA, should be to obtain sample products, that is, plans. A questionnaire should inquire about the characteristics of an optimum plan and an unsuccessful one.

Having acquired a first approximation of the elements with which an expert works, the knowledge acquirer should then try to capture a description of what the planner actually does at the action level. This is not always a straightforward process. What an expert says he does, and what he actually does, may differ. Equally adept experts may solve the same problem differently by using different decision aids, for instance.

METHODS

Some methods are better than others at capturing particular dimensions of the problem-solving activity. The knowledge acquirer's learning curve in the domain may make one method more useful than another, and the expert may respond better to one method than to another. Consequently, two and, if possible, three methods should be used to acquire the knowledge necessary to produce a description of the planning activity. It is anticipated that a good deal of knowledge will be acquired by using a number of methods, especially interviews. But the major unifying goal of knowledge acquisition at this point should be to capture the knowledge required to produce a task description (or facsimile)—the process involved in producing a plan. In other words, the steps in the process should be described in terms of verbs that describe the activity.

First, interviews with an expert should be conducted. A walk-through, using a simple scenario as a prompt, should follow. (In the case of the ASPA project, conducting postplanning interviews with airwing officers who had recently produced a subdomain plan in Naval Strike Warfare Center's classes is yet another possibility.) Follow-on interviews, designed to fill in holes in the task description, are also anticipated.

After the interviews and walk-throughs, efforts should focus on those parts of the planning process that involve "calculating," "analyzing," "determining," or other cognitive subgoal tasks. Words like these provide place markers that point to places where more in-depth knowledge acquisition is required.

Answers to questions about *how* something is calculated, analyzed, or determined point to heuristic rules and documents used by the expert. The answers should also produce information that can be reformulated as subgoal "inputs" and "outputs."

A summary chart of recommended knowledge acquisition methods, and anticipated results, is given in Table 1. Examples are given to show the reader the kinds of knowledge the methods are believed to produce, in two ways: "Product" refers to the specific type of knowledge acquired from the expert; "Knowledge structure" refers to the category of the knowledge acquirer's understanding to which the type of knowledge will contribute. The order roughly corresponds to ideas about how the knowledge acquirer can most efficiently capture, analyze, interpret and refine subdomain knowledge.

Table 1. Knowledge acquisition methods.

Knowledge Acquisition Phase & Method	Knowledge source	Product	Knowledge structure	Examples
<u>First Phase</u>				
Textbook review	Expert referral	Problem definition	Overview description	"A successful weaponeering plan achieves the desired level of target destruction with the minimum use of resources (weapons, aircraft) compatible with the least jeopardy to assets (aviators, aircraft and equipment)"
Analysis of available Information	Expert interviews & expert documents	Problem definition	Overview description	
Questionnaire	Subdomain experts	Problem definition	Overview description	
Identification of expert(s)				
<u>Second Phase</u>				
Compilation of available information	Expert documents	Glossary	Names of classes of objects, acronyms, etc.	PK = probability of kill LGBs = laser-guided bombs GBU = guided bomb units FMPL = fleet mission program library
Analysis & compilation of available information	Expert interviews & expert documents	Table(s) of resource elements	Expert's basic factual knowledge	TOT = time on target ROI = rules of engagement
Read through with expert(s)	Expert documents	Table(s) of resource elements	Expert's basic factual knowledge	Names of: - Weapons - Fuze - FINS - Other equipment
Interview	Subdomain expert(s)	Complete, reviewed resource table(s)	Expert's basic factual knowledge	

Table 1. Knowledge acquisition methods (continued).

Knowledge Acquisition Phase & Method	Knowledge source	Product	Knowledge structure	Examples
<u>Third Phase</u> Questionnaire	Subdomain expert(s)	Definition of essential, optimum, & typical information	Refined problem definition	<p>Target description Target-area threats Target-area environmental conditions (terrain, weather, ceiling, visibility)</p> <p>Quantities and kinds of bombs, fuzes, functioning-delay times, arming times, fins, required airspeed, delivery parameters, etc.</p>
<u>Fourth Phase</u> Interviews	Subdomain expert(s)	Sample planning products	Refined problem definition	WEAPON SELECTED with highest probability of achieving desired level of damage and which incurs minimal risk
Interview/questionnaire	Subdomain expert(s)	Definition of optimum and suboptimum plans	Refined problem definition (identify planning goals)	vs. wrong weapon, no damage, reoccurring risk by having to return
<u>Fifth Phase</u> Interviews	Subdomain expert(s)	Description of planning processes	Reasoning sample process steps	"First, I look at the target, and try to find a weapon with a PK as close to 1 as possible. Then, I look at target-area tactics"
Walk through	Subdomain expert(s)/ simple scenario	Description of planning processes	Reasoning sample process steps	a. Analyzes target b. Goes to JMEM's Visual Delivery
Reinterview	Subdomain expert	Description of planning processes	Reasoning sample process steps	Manual to identify a weapon with a high probability of kill c. Analyzes threat d. Analyzes environment

Table 1. Knowledge acquisition methods (continued).

Knowledge Acquisition Phase & Method	Knowledge source	Product	Knowledge structure	Examples
Interviews	Task description/expert(s)	Detailed description of expert's procedure for accomplishing subgoal tasks (algorithms, principles, rules, and tables)	Expert's reasoning (inputs and outputs of subgoal problem solution)	<p>Threat analysis:</p> <ul style="list-style-type: none"> a. Identify kind of threat (e.g., early detection, fighters, AAA, SAMs) b. Identify threat envelopes <p>Procedure:</p> <ul style="list-style-type: none"> a. Generate plan to suppress, neutralize as many threats as possible b. Identify speed, maneuverability, altitude requirements in light of threat c. Look up weight and drag counts in TAC Manual tables
Walk throughs	Task description/expert(s)	Detailed description of computer-aided subgoal problem solving	Expert's reasoning (inputs and outputs of subgoal problem solution)	
Read throughs	Task description/expert(s)	Detailed description of documentary data/tables used to solve subgoal problems	Expert's reasoning (inputs and outputs of subgoal problem solution)	<p>Rules:</p> <ul style="list-style-type: none"> a. AAA cannot be suppressed. b. Reject weapons where delivery parameters require aircraft to fly in heart of threat envelope. c. Reject weapons loads where weight and drag counts impair fuel, minimum speed, and maneuverability requirements.

REFERENCES

- Bainbridge, L. (1981). Verbal Reports as Evidence of the Process Operator's Knowledge. *Fuzzy Reasoning and its Applications*. E.H. Mandani and B.R. Gaines (Eds.). London, England: Academic Press, 343-368.
- Barr, A. and E.A. Feigenbaum (Eds.) (1982). *The Handbook of Artificial Intelligence*. Stanford, California: Heuris Tech Press.
- Barstow, D.R. (1979). *Knowledge Based Program Construction*. New York: Elsevier Science Publishing Co., Inc.
- Blau, P.M. (1963). *The Dynamics of Bureaucracy*. Chicago: Chicago University Press.
- Brandt, R.M. (1972). *Studying Behavior in Natural Settings*. New York: Holt, Rhinehart & Winston.
- Brayfield, A.H. and W.H. Crockett (1955). Employee Attitudes and Employee Performance. *Psychological Bulletin*, 52, 396-424.
- Brooks, H. (1983). *Information Retrieval and Expert Systems: Designing a Knowledge Base*. Presentation at Informatics 7. Conference held by the ASLIB Informatics Groups and the British Computer Society Information Retrieval Group, Penbroke College, Cambridge, England, March 1983.
- Buchanan, B.G. (1979). Issues of Representation in Conveying the Scope and Limitations of Intelligent Assistant Programs. *Machine Intelligence*, Vol. 9. J.E. Hayes, D.M. Michie, and L.I. Mikulich (Eds.). Chichester, England: Ellis Horwood, Ltd., 407-425.
- Chandrasekaran, B. (1985). Generic Tasks In Expert System Design and Their Role in Explanation of Problem Solving. Invited presentation at the National Academy of Science/Office of Naval Research Workshop on AI and Distributed Problem Solving, May 16-17, 1985.
- Cicourel, A. (1964). *Method and Measurement in Sociology*. New York: The Free Press.
- Clancey, W.J. (1981). Methodology for Building an Intelligent Tutoring System. Computer Science Dept. Report No. STAN-CS-81-894, Stanford University, Palo Alto, California.
- Dalton, M. (1959). *Men Who Manage*. New York: Wylie.
- Dalton, M. (1967). Preconceptions and Methods in Men Who Manage. *Sociologists at Work*. P.E. Hammond (Ed.). New York: Anchor Books.

- deGreef, P. and J. Breuker. (1985). A Case Study in Structured Knowledge Acquisition. *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, 390-392.
- Ehrler, P., G. Osga, D. Lambert, and T. Cook. (1985). Expert Systems in Support of Composite Warfare Commander Mission Planning. NOSC Technical Report 1067. Naval Oceans Systems Center, San Diego, California.
- Ericsson, K.A. and H.A. Simon. (1984). *Protocol Analysis*. Massachusetts: MIT Press.
- Ericsson, K.A. and H.A. Simon. (1980). Verbal Reports as Data, *Psychological Review*, 87(3), 215-251.
- Feigenbaum, E.A. and P. McCorduck. (1983). *The Fifth Generation*. New York: Addison-Wesley.
- Flammer, A. (1981). Towards a Theory of Question Asking. *Psychological Research*, 43, 407-420.
- Flammer, A., A. Grob, T. Leuthardt, and R. Luthi. (1984). Asking How to Act. *Archives de Psychologie*, 52, 103-120.
- Flammer, A., H. Kaiser, and P. Mueller-Bouquet. (1981). Predicting What Questions People Will Ask. *Psychological Research*, 43, 421-429.
- Fox, M.S. (1983). Techniques for Sensor-Based Diagnosis. Proceeding of 8th International Joint Conference on Artificial Intelligence.
- Frankel, R.M. (1982). Impact: An Interactional-Based Method for Preserving and Analyzing Clinical Transactions. *Exploration in Provider Patient Interactions*. Pettigrew (Ed.). Chicago, Illinois: Humana Press.
- Garfinkel, H. (1969). *Studies in Ethnomethodology*. New Jersey: Prentice Hall.
- Gouldener, A.W. (1954). *Patterns of Industrial Bureaucracy*. New York: The Free Press.
- Grover, M.D. (1983). Pragmatic Knowledge Acquisition Methodology. Proceedings of the 8th International Joint Conference on Artificial Intelligence.
- Harrison, A. and R. Bramson. (1982). *The Art of Thinking*. California: Anchor Press.
- Hartley, R.T. (1982). A Conceptual Basis for Expert Systems Methodology. Proceedings of the 2nd British Computer Society Conference on Expert Systems, Brunel University, Surrey, England, 47-51.

- Hyman, H.H., W.J. Cobb, J.J. Feldman, G.W. Hart, and C.H. Slembert (1954). *Interviewing in Social Research*. Chicago, Illinois: University of Chicago Press.
- Kinsey, A.C., W.B. Pomeroy, C.E. Martin, and P.H. Gebhard, P.H. (1965). *Sexual Behavior in the Human Female*. New York: Pocket Books.
- Lehner, P.E., D. Zirk, R.B. Hall, and L. Adelman. (1985). Human Factors in Rule Based Systems. Par Technology Corporation. Report Number 85-l09. McClean, Virginia.
- Mittal, S. and C.L. Dym. (1985). Knowledge Acquisition From Multiple Experts. *The AI Magazine*, Summer, 32-36.
- Miyake, N. and D. Norman. (1979). To Ask a Question, One Must Know Enough to Know What Is Not Known. *Journal of Verbal Learning and Verbal Behavior*, 18, 357-364.
- Natale, M., E. Elliot, and J. Jaffe. (1979). Vocal Interruptions in Dyadic Communication as a Function of Speech and Social Anxiety. *Journal of Personality and Social Psychology*, 37, 865-78.
- Rogers, C.R. (1959). A Theory of Therapy, Personality, and Interpersonal Relationships, as Developed in the Client Centered Framework. *Psychology: A Study of a Science*. S. Koch (Ed.). Vol III. New York: McGraw-Hill.
- Rogers, W. and S. Jones. (1975). Effects of Dominance Tendencies on Floor Holding and Interruption Behavior in Dyadic Interaction. *Human Communication Research*, 1, 113-22.
- Sacks, H., E. Schegloff, and G. Jefferson. (1974). A Simplest Systematics for the Organization of Turn-Taking for Conversation. *Language*, 50, 696-735.
- Selznick, P. (1949). *TVA and the Grassroots*. Berkeley: University of California Press.
- Smith, R.G. and J.D. Baker. (1983). The Dipmeter Advisory System: A Case Study in Commercial Expert System Development. Proceedings of 8th International Joint Conference on Artificial Intelligence, 122-129.
- Sowa, J.F. (1984). *Conceptual Structures: Information Processing in Mind and Machine*. Massachusetts: Addison-Wesley Publishing Company.
- Thompson, J.D. and N.J. Demerath. (1952). Some Experience with Group Interviews. *Social Forces*, 31, 148-154.
- Waterman, D.A. and B. Jenkins. (1979). Heuristic Modelling Using Rule-Based Computer Systems, *Terrorism: Threat, Reality, Response*. R. Kupperman and D. Trent (Eds.). California: Hoover Institution Press.

- Weick, K.E. (1969). Meaning and Misunderstanding. Review of H. Garfinkel: Studies in Ethnomethodology. *Contemporary Psychology*, 14, 357-360.
- Welbank, M. (1983). A Review of Knowledge Acquisition Techniques for Expert Systems. Martlesham Consultancy Services, British Telecom Research Laboratories, Martlesham Heath, Ipswich, England.
- Wernimont, P.F., and J.P. Campbell. (1968). Signs, samples and criteria. *Journal of Applied Psychology*, 52, 372-376.
- Wicker, A.W. (1969). Attitude versus action: The relationship of Verbal and Overt Behavioral Responses to Attitude Objects. *Journal of Social Issues*, 25, 41-78.
- Zimmerman, D. and C. West. (1975). Sex Roles, Interruptions and Silences in Conversation. *Language and Sex: Difference and Dominance*. B. Thorne and Nancy Henley (Eds.). Massachusetts: Newbury House, 105-29.

BIBLIOGRAPHY

- Bennett, J.S. (198?). On the Structure of the Acquisition Process for Rule Based Systems. In *Machine Intelligence: Infotech State of the Art Report*, Series 9, No. 3, Pergamon Infotech, Ltd., Maidenhead, Berks, England, 189-204.
- Buchanan, B.G., D. Barstow, R. Bechtel, J. Bennett, W. Clancey, C. Kulikowski, T. Mitchel, and D.A. Waterman. (1983). Constructing an Expert System. *Building Expert Systems*. F. Hayes-Roth, D.A. Waterman, and D.B. Lenat (Eds.). Reading: Addison-Wesley, 127-168.
- Chandrasekaran, B. (1983). Expert Systems: Matching Techniques to Tasks. Defense Technical Information Report, AFOSR-TR-83-0660, Alexandria, Virginia.
- Charniak, E. (1985). *Introduction to Artificial Intelligence*. Massachusetts: Addison-Wesley Publishing Co.
- Cicourel, A. (1982). Language and Belief in a Medical Setting. *Contemporary Perceptions of Language: Interdisciplinary Dimensions*. H. Byrnes (Ed.). Washington D.C.: Georgetown University Press.
- Feigenbaum, E.A. (1979). Themes and Case Studies of Knowledge Engineering. *Expert Systems in the Microelectronic Age*. D. Michie (Ed.). Edinburgh, Scotland: Edinburgh University Press, 3-25.
- Frankel, R. M. (1984). From Sentence to Sequence: Understanding the Medical Encounter through Microinteractional Analysis. *Discourse Processes*, 7, 135-170.
- Gorry, G. A. (1973). Computer-Aided Clinical Decision-Making. *Methods of Information in Medicine*, 12, 1, 45-51.
- Hillman, D.J. (1985). Artificial intelligence. *Human Factors*, 27(1), 21-31.
- Kornhauser, A. and P. Sheatsley. (1959). Questionnaire Construction and Interview Procedure. *Research Methods in Social Relations*, C. Sellitz, M. Johada, M. Deutsch and S. Cook (Eds.), New York: Holt, Rinehart and Winston.
- Sutcliffe, J.P. (1972). On the Role of "Instructions to the Subject" in Psychological Experiments. *American Psychologist*, 27, 755-760.

END

9-87

DTIC